

What type of transparency in OTC markets?*

Maren Vairo[†] and Piotr Dworczak[‡]

August 11, 2023

Abstract

Financial over-the-counter markets have been traditionally very opaque. Recent regulation promotes transparency in some of these markets by lowering search costs, allowing traders to request quotes from multiple dealers at the same time (pre-trade transparency), and requiring public disclosure of past transactions (post-trade transparency). We evaluate these policies using a dynamic trading model with adverse selection. We show that post-trade transparency improves upon the opaque market but is dominated by pre-trade transparency; moreover, adding post-trade transparency to a pre-trade transparent market offers no benefits and can be harmful. We identify cases in which lowering search costs can be detrimental to market efficiency. Finally, relying on a mechanism-design approach, we characterize the optimal trading mechanism.

Keywords: OTC markets, pre-trade transparency, post-trade transparency, search, price discovery

*We thank Belén Jerez, Josh Mollner, Humberto Moreira, Alessandro Pavan, Harry Pei, Michael Peters, and various seminar audiences for helpful comments and suggestions. Dworczak gratefully acknowledges the support of the Alfred P. Sloan Fellowship.

[†]Department of Economics, Northwestern University.

[‡]Department of Economics, Northwestern University.

1 Introduction

Many important financial assets, including corporate bonds, swaps, currencies, non-standard derivatives, real estate, and consumer products, are primarily traded in over-the-counter (OTC) markets. *Opacity* has been traditionally one of the key features setting OTC markets apart from traditional centralized exchanges (Duffie, 2012). Traders in financial over-the-counter markets have to search for counterparties before they transact (Duffie et al., 2005). Often, they must request a quote from an OTC dealer and decide whether to accept it before they can learn what prices other dealers may be offering (Zhu, 2012). This unawareness of the currently available quotes can be described as lack of *pre-trade transparency*. Similarly, in most OTC markets, traders do not observe each others' transactions, and hence cannot rely on data about past prices to inform their trading strategies or refine their estimates of asset values. This feature is generally referred to as lack of *post-trade transparency*.

The opacity of financial OTC markets has been raising concerns among policymakers, which were exacerbated by 2007-2008 financial crisis. In the US, a number of regulations increased both the pre- and post-trade transparency of some OTC markets. Already in early 2000s, the SEC instituted the Trade Reporting and Compliance Engine (TRACE) in selected segments of the corporate bond market, with the set of TRACE-eligible securities expanding over time.¹ Under TRACE, broker-dealers who are FINRA members are required to disseminate price and volume information about their trades, effectively in real time. This results in considerably increased post-trade transparency in markets covered by TRACE. By providing access to data about past prices, the regulator intended to promote market efficiency by “leveling the playing field.” More recently, in the aftermath of the so-called Libor scandal, regulators around the world took actions to support robust financial benchmarks (e.g., the replacement of Libor with SOFR in the US) that are also an important source of post-trade transparency through public dissemination of *average* prices.² The 2010 Dodd-Frank act, on the other hand, promoted pre-trade transparency in swap markets by introducing Swap Execution Facilities (SEFs). In contrast to traditional search markets, SEFs were designed to allow investors to request quotes from multiple dealers at the same time.

This paper addresses the problem of *optimal* design of transparency in a financial over-the-counter market. We propose a simple dynamic model of trading under adverse selection in which asymmetrically informed dealers compete to trade with (initially) uninformed investors. We investigate how different forms of pre- and post-trade transparency affect the equilibrium level of welfare. In the first part of the paper, we compare the performance of

¹See, for example, Asquith et al. (2019).

²See, among others, Hou and Skeie (2013), Duffie and Stein (2015), and Duffie et al. (2017).

the market under several regimes created by the regulation discussed above. In the second part, we use mechanism design to identify the optimal market design in a broad class of dynamic trading mechanisms. We view the main contribution of the paper as providing the first theoretical framework that accommodates both pre- and post-trade transparency, resulting in novel insights that could guide future regulation of OTC markets.

Our model features two OTC dealers,³ who are long-run players competing for the order flow, and an infinite sequence of investors who come to the market desiring to buy or sell a unit of the asset. The asset has a common value component but the investor's value differs from the common value by a constant private-value component resulting, for example, from a liquidity need. Thus, there is a constant, strictly positive gain from trade between investors and dealers, and full efficiency requires trade between a dealer and an investor in every period. An adverse selection problem precludes this efficient outcome: Dealers observe private signals about the common value of the asset, while investors are initially uninformed about the asset value.

We ask how different transparency regulations affect the equilibrium probability of trade (which is equivalent to a standard measure of efficiency in our framework). We model the traditionally opaque OTC market as a search market in which an investor must search sequentially (subject to a cost) to obtain quotes from dealers. The investor does not have access to any information about past trades, and can only find out about dealers' quotes by visiting them. The search cost acts as a simple proxy for the market pre-trade opacity (it is the search cost that prevents investors from fully learning about the currently available prices). Thus, we can think of lowering the search cost (perhaps due to regulation requiring electronic trading) as an increase in market transparency. We model the fully pre-trade transparent market as a first-price auction between the dealers, motivated by the popular request-for-quote (RFQ) protocol featured on many trading platforms. Finally, we model post-trade transparency by letting investors who come to the market observe the history of transactions. A fully transparent market has both pre- and post-trade transparency.

Our framework endogenizes the influence that these different transparency regimes have on equilibrium learning and information asymmetry. Investors are initially uninformed but they can learn about the asset value from dealers' quotes in three ways, depending on the setting: *(i)* by searching and requesting quotes (in opaque regimes), *(ii)* by observing dealers' quotes in the auction (in pre-trade transparent regimes), and *(iii)* by inspecting dealers' past executed quotes (in post-trade transparent regimes). We show that these three information channels differ significantly in terms of their equilibrium consequences for dealers' pricing strategies, and hence welfare. In opaque markets, dealers' informational advantage creates an

³In Section 5, we comment on how our results can be extended to any finite number of dealers.

adverse selection problem for the investor, and supporting an equilibrium requires quotes that are often rejected.⁴ In pre-trade transparent markets, direct competition exposes the dealers to the adverse selection problem, and the least informed of them must provide conservative quotes to break even. Finally, in post-trade transparent markets, dealers have incentives to manipulate future beliefs about the asset value with their quotes, and this signaling motive imposes constraints on equilibrium trading probabilities.

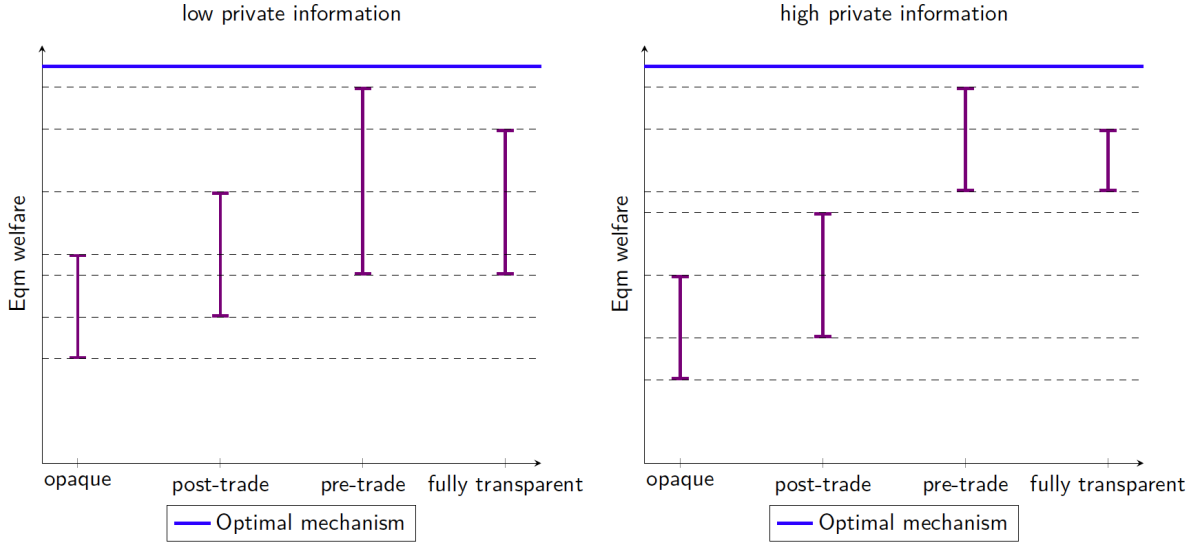
Figures 1.1 and 1.2 summarize our findings in the first part of the paper. Unsurprisingly—given the presence of signaling motives and adverse selection—our model features equilibrium multiplicity. We impose only a few intuitive equilibrium refinements motivated by practical considerations, and compare different regimes using two dominance notions that do not rely on equilibrium selection. Specifically, one market regime is said to *weakly welfare-dominate* another one if it has higher expected welfare in both the best and the worst equilibrium; *strong welfare dominance* requires that the worst equilibrium of one regime has higher expected welfare than the best equilibrium of another regime. Welfare comparisons turn out to depend on the value of a parameter that governs the informational advantage of informed dealers (labeled “low” and “high” private information in Figures 1.1 and 1.2). The parameter is higher in markets where dealers are more likely to possess private information relevant to the asset value.

We find that post-trade transparency improves upon the opaque market regardless of the level of search costs; however, the improvement is only in the weak sense. Pre-trade transparency welfare-dominates post-trade transparency, and the improvement is in the strong sense in markets with a high degree of private information. Perhaps surprisingly, adding post-trade transparency to a pre-trade transparent market can only hurt welfare: The worst equilibrium stays the same, while the best equilibrium becomes less efficient.

A high-level intuition for these results is as follows. Post-trade transparent markets become efficient once past quotes reveal the value of the asset. However, due to signaling incentives of dealers, they can be quite inefficient when the value of the asset is uncertain. Pre-trade transparency is highly efficient for two reasons: First, dealers’ information becomes symmetric after they participate in an auction; second, direct competition between dealers makes them offer attractive quotes. The detrimental effect of adding post-trade transparency is explained by the insight that once dealers are symmetrically informed and compete against each other, investors do not *need* to be informed to make efficient trading decisions. An attempt to make investors more informed by disclosing past transactions “backfires” as it introduces a temptation for the dealers to manipulate future investors’ beliefs.

⁴An efficient pooling equilibrium is ruled out in our analysis by assuming that private gains from trade are not too large compared to the uncertainty about the common value component.

Figure 1.1: Welfare ranking across trading regimes




We also find an ambiguous effect of lowering search costs: It tends to improve welfare in markets where the informational advantage of dealers is not too large, but it leads to very inefficient outcomes in markets with a lot of information asymmetry. The reason is that less informed dealers face a strong form of adverse selection in such markets: A large fraction of their order flow comes from investors who are shopping around for a better quote after seeing a quote of an informed dealer. As a consequence, these less informed dealers essentially refuse to trade in equilibrium, as they would make losses otherwise.⁵

This result stays in sharp contrast to our previous finding about the attractiveness of competition between dealers in the pre-trade transparent market. It would be natural to conjecture that the search market approaches the same highly efficient outcome as search costs vanish; yet, the opposite can be true, with the search market approaching the lowest possible level of welfare in the unique equilibrium. This is because the details of the protocol matter for the effects of competition in the face of adverse selection. Our framework highlights two important features of pre-trade transparency that even very competitive search markets lack: Simultaneity of quotes (which eliminates the adverse inference the uninformed dealer makes from the event of being asked for a quote) and order visibility (which reduces the information asymmetry across dealers at subsequent trading opportunities).

While we cover a number of different regimes in our equilibrium analysis, it is impossible

⁵This mechanism is reminiscent of the adverse selection problem faced by the market maker in classical models such as Kyle (1985) and Glosten and Milgrom (1985) but the difference is that adverse selection in our model is endogenously exacerbated by the equilibrium search behavior when search costs are low.

Figure 1.2: The effects of lowering search costs (s)



to analyze all trading mechanisms used in practice. Regulation leaves a lot of scope for market participants to experiment with the designs of trading mechanisms. For example, in the US swap market, traders can choose from several dozen SEFs that differ in their pre- and post-trade transparency, as well as other features (see [SIFMA, 2016](#)). In the second part of the paper, we turn to the problem of *optimal* market design. We characterize the welfare-maximizing dynamic mechanism subject to natural constraints, such as market-clearing, budget-balance, incentive-compatibility and individual rationality.⁶ We also require the indirect implementation to be state- and time-invariant, since the rules of real-life trading mechanisms are fixed and rarely depend on the market circumstances.

We find that the first-best outcome (trade with probability one) can be robustly implemented for a wide range of parameter values. When the first best cannot be implemented, the optimal mechanism features interior probability of trade in case both dealers are uninformed about the value of the asset. Intuitively, in such cases, if dealers traded with probability one even conditional on receiving uninformative signals, some informed dealers would be able to secure a better price by behaving as if they were uninformed.

To gain intuition for the properties of the optimal mechanism, we construct an indirect trading mechanism that implements the optimal outcome in the best-case equilibrium. The mechanism is static, and resembles an open ascending or descending auction (depending on whether the investor is a buyer or a seller) combined with an ex-post bargaining stage.

⁶We impose an ex-post version of incentive constraints to circumvent the well-known issue with mechanism design models with correlated information, as explained in [Cr mer and McLean \(1985\)](#).

The mechanism features no post-trade transparency, that is, it discloses no information to traders who did not participate in the auction. Intuitively, an open second-price-like auction maximizes the amount of learning and alleviates the adverse selection problem for the participants. Because dealers observe the quotes that their competitors were willing to offer, they acquire symmetric information, which results in efficient trading in subsequent periods, without the need to broadcast information to future investors. Thus, the optimal mechanism confirms the intuition from the first part of the paper: It is crucial to ensure symmetric information and direct competition between the dealers, while attempts to disclose information to investors are futile at best and detrimental for welfare at worst. Finally, the post-auction bargaining stage corrects a potential inefficiency associated with the fact that the investor—say, a buyer—trades at the second-lowest quote in a descending auction. It is possible that the buyer updates her estimate of the value of the asset and would like to buy at a price that lies between the second-lowest quote and the value of the asset to the winning dealer. In this case, the post-auction bargaining stage allows the dealer and the investor to agree on a price that makes trade profitable for both. This last finding explains why it may be efficient for platforms to allow for bilateral negotiations on top of offering an auction-like protocol; many, though not all, SEFs offer such a feature.⁷

The rest of the paper is organized as follows. We review the related literature next. In Section 2, we introduce the model. In Section 3, we present our main results on welfare properties of equilibria under different transparency regimes. In Section 4, we study the optimal mechanism. In Section 5, we discuss extensions and limitations of the model, focusing on how they shape policy implications of our results. All proofs are collected in the appendix.

1.1 Related literature

There is a sizable theoretical literature devoted to the analysis of transparency of financial OTC markets. Our primary contribution to this literature is that we propose a single modeling framework that allows for a tractable analysis of both pre- and post-trade transparency (as well as the effects of lowering search costs, and varying other aspects of the market design). This leads to novel insights—for example, about the interaction of the two types of transparency—that would not be possible to obtain otherwise. To the best of our knowledge, all previous studies focused on evaluating one aspect of transparency at a time.

Similarly to us, [Asriyan et al. \(2017\)](#) analyze a dynamic model with adverse selection. They investigate the effects of post-trade transparency and find equilibrium multiplicity and an ambiguous effect of disclosure of past trades on welfare. However, their results are driven by a very different mechanism: They focus on spillover effects across markets, and the

⁷MarketAxess, DelphX, TradeWeb, and Bloomberg platforms serve as an example; see [SIFMA \(2016\)](#).

fact that sellers may strategically delay their trades in anticipation of information coming from markets for correlated assets. [Duffie et al. \(2017\)](#) analyze the transparency effects of introducing a benchmark to an opaque search market, similar to our opaque market (our way of modeling search is even closer to [Zhu, 2012](#)). They show that benchmarks raise surplus for a large set of parameters. We analyze benchmarks briefly as an extension of our post-trade transparent regime and argue that welfare consequences become much more subtle when *(i)* there is adverse selection, and *(ii)* the benchmark is computed endogenously from dealers' quotes, as opposed to assumed exogenous (as in [Duffie et al., 2017](#)). [Ollar et al. \(2021\)](#) analyze the effects of disclosure of past trades in a dynamic uniform-price double auction—they find that equilibrium existence puts restrictions on how transparent the market can be. While we study a very different model, we also find that separating equilibria could fail to exist when dealers can benefit from manipulating beliefs in the market. [Kakhbod and Song \(2020\)](#) (see also [Kakhbod and Song, 2022](#)) analyze a model that is similar to ours and find that post-trade transparency can have an adverse effect on price informativeness: In a repeated interaction between traders, disclosure of past prices allows for richer punishments for off-path behavior, which in turn makes it possible to support opaque pooling equilibria for a larger set of parameters. Thus, paradoxically, more post-trade transparency can lead to less information being disclosed on equilibrium path.⁸ Our paper is complementary as it studies a different objective (allocative rather than price efficiency) and focuses on separating equilibria (we rule out pooling equilibria with an assumption about model parameters). An argument in favor of post-trade transparency can be found in [Back et al. \(2020\)](#) who show that dealers may offer more attractive prices to sellers of assets when these prices are observed by prospective buyers. However, in their model, the benefits of signaling are strengthened by the assumption that trade always happens between a dealer and a buyer (effectively, the dealer has an infinite cost for holding the asset). Finally, [Glebkin et al. \(2023\)](#) look at the effects of introducing pre-trade transparency (an auction between multiple traders) in the canonical OTC model of [Duffie et al. \(2005\)](#). Their analysis does not feature adverse selection and focuses on the equilibrium distribution of asset holdings in a large market.

On the empirical side, a series of papers have studied the effects of increased post-trade transparency, through the introduction of TRACE, in the US corporate bonds market ([Edwards et al., 2007](#), [Bessembinder and Maxwell, 2008](#), [Asquith et al., 2019](#)). The findings in these papers indicate that post-trade transparency led to a reduction in transaction costs. However, [Asquith et al. \(2019\)](#) find that this reduction did not necessarily translate into

⁸[Kakhbod and Song \(2020\)](#) show that this increased opacity may, however, be beneficial for market liquidity. A similar force is uncovered by [Bergemann and Hörner \(2018\)](#) in the context of a repeated first-price auction.

higher trading activity. They also document that the effect of TRACE on market liquidity differs considerably, both in sign and in magnitude, across bond types. [Aquilina et al. \(2022\)](#) find that increasing the informativeness and precision of benchmarks in the swaps market led to improvements in the quality of dealer-trader matches, thereby increasing the volume of welfare-enhancing trades. Overall, these empirical findings are consistent with the results in this paper which state that the effects of post-trade transparency on market welfare may be ambiguous and depend on the details of the market (and possibly on equilibrium selection as well). With regard to pre-trade transparency, [Hendershott and Madhavan \(2015\)](#) analyze the effects of electronic trading, through periodic first-price auctions, on the corporate bond market. They document that the resulting reduction in search costs and increase in competition across dealers led to more competitive prices and higher market liquidity, which is also in line with our findings. Like us, [Riggs et al. \(2020\)](#) find that market imperfections, such as private information and existing relationships between dealers and customers, may prevent pre-trade transparent regimes from being fully competitive. Using SEF trading data, they document that customers contact only a few dealers, and that some dealers refrain from responding to requests. They microfound these findings by proposing a model where dealers are exposed to a form of winner’s curse that—unlike in our paper—manifests in dealers’ future (in)ability to offload any acquired positions in the inter-dealer market. Our paper highlights complementary forces (such as adverse selection stemming from common values) that may drive market failure in auction-type markets.

Since our framework features various trading protocols—including search, bargaining, and auctions—it inevitably connects to many classical literatures on adverse selection, in particular in the context of search (e.g., [Wolinsky, 1990](#), [Blouin and Serrano, 2001](#), [Lauer-
mann and Wolinsky, 2016](#)) and bargaining (e.g., [Hörner and Vieille, 2009](#), [Fuchs et al., 2016](#), [Kim, 2017](#), [Kaya and Kim, 2018](#)). These literatures studied the influence of information disclosure on allocative efficiency and information aggregation through prices. While certain forces that drive our welfare results are familiar from these models, to the best of our knowledge, pre- and post-trade transparency have not been studied together within a single framework. In [Appendix A](#), we review these literatures in more detail and explain both the similarities and differences to our setting.

2 Model

2.1 The trading environment

We consider a simple trading environment with discrete time and infinite horizon, $t = 0, 1, \dots$, a single risky asset, two long-lived dealers, and a sequence of short-lived customers.

Asset. The risky asset has a common value $v_t \in \{-1, 1\}$ at time t . At time $t = 0$, v_0 is drawn from the uniform distribution. In every subsequent period t , the value of the asset is either “persistent” ($v_t = v_{t-1}$) or the value of the asset “resets,” that is, it is redrawn from the uniform distribution. We interpret the reset of the value of the asset as arrival of exogenous payoff-relevant news to the market.⁹

Some persistence in the asset value is necessary to analyze the effects of transparency but persistence complicates the analysis in an infinite-horizon model with learning: We model the evolution of the value of the asset in a stylized and parsimonious way by assuming that the asset value can be persistent over at most two periods. (In Section 5, we comment on why we believe our insights would continue to hold qualitatively without that assumption.) Specifically, there is an underlying Markov chain $Z_t \in \{0, 1\}$, with $Z_t = 0$ representing the event that v_t gets redrawn in period t (news arrives to the market), and $Z_t = 1$ representing the event that $v_t = v_{t-1}$ (no news arrive to the market). Conditional on news arriving to the market at time $t - 1$, there is probability $\gamma \in (0, 1)$ that no news will arrive in period t . Formally, $\Pr(Z_t = 1|Z_{t-1} = 0) = \gamma$. However, news arrives for sure after a period with no news—i.e., $\Pr(Z_t = 1|Z_{t-1} = 1) = 0$. Thus, γ is a measure of persistence of the asset’s value over time. Under stationarity assumptions, equilibrium behavior needs to be characterized only conditional on one of the two possible states—the “reset state” $Z_t = 0$ and the “persistent state” $Z_t = 1$ —which greatly simplifies the analysis.¹⁰

Dealers. Two dealers, labeled $i \in \{A, B\}$, compete in the market for the risky asset by providing quotes to customers, according to a trading protocol that will be specified later. Both dealers value each unit of the asset at v_t at time t , and face no inventory or short-selling constraints. They discount the future through a common discount factor $\delta \in (0, 1]$. Formally, if in each period t a dealer sells $x_t^s \in \mathbf{R}_+$ units of the asset at a price a_t and buys $x_t^b \in \mathbf{R}_+$ units at a price b_t , her intertemporal payoff is

$$\sum_{t=0}^{\infty} \delta^t (x_t^s (a_t - v_t) + x_t^b (-b_t + v_t)).$$

Dealers differ in the information they have about the asset value. Conditional on news arriving to the market in period t ($Z_t = 0$), each dealer observes a signal realization at the beginning of period t . Conditional on v_t , the two signals are independent, and reveal the true state v_t with conditional probability $q \in (0, 1)$; with the remaining probability $1 - q$,

⁹Nothing changes in our analysis if we interpret v_t as the *change* in the asset’s value, so that the value follows a random walk, as long as previous values become common knowledge when the value resets.

¹⁰While we find the infinite-horizon version of our model more natural to interpret, we note that all results would continue to hold in a simpler two-period model with a fully persistent state but with an exogenous probability $1 - \gamma$ that the game ends (e.g., the asset is liquidated) after the first period.

the signal conveys no information about v_t , leading to a posterior belief equal to the prior. If no news arrive to the market in period t ($Z_t = 1$), dealers do not observe any new signal realizations at t . Signal realizations determine dealers’ types. We will refer to a dealer who observed the uninformative signal realization as an “uninformed type;” we will refer to a dealer who observed the revealing signal realization as an “informed type;” we will also use the “high type” (resp., “low type”) to refer to an informed dealer who knows that $v_t = 1$ (resp., $v_t = -1$). We will index these three types by $\sigma \in \{l, u, h\}$.

Our assumptions about the information structure imply that while news arrival to the market (state Z_t) is common knowledge, the dealers may not succeed in “deciphering” the implications of the news for the asset value. Alternatively, as in the models of [Budish et al. \(2015\)](#) and [Baldauf and Mollner \(2020\)](#), dealers may experience random latency between the arrival of news to the market and the time this information is reflected in their quotes. A key feature of the information structure is that a dealer may be less informed than her competitor in some periods.¹¹

Customers. In each period t , a short-lived customer (investor) arrives to the market, desiring to trade one unit of the asset. The customer is a buyer or a seller with equal probability. Assuming that the equilibrium is symmetric across these two scenarios, we can without loss of generality focus on the case when the customer wants to buy the asset.¹² The value for the period- t customer is then $v_t + \Delta$, where $\Delta \in (0, 1)$ is a known constant, reflecting an immediacy need or a similar buyer-specific benefit for obtaining the asset. Therefore, the payoff of a period- t customer-buyer who trades $x \in \{0, 1\}$ units at a price a is $x(-a + v_t + \Delta)$, while the payoff of a customer-seller trading $x \in \{0, 1\}$ units at a price b is $x(b - v_t - \Delta)$.

Customers possess no private information but they may be able to observe some public signals revealed by the market protocol. (In [Appendix OA.1](#), we extend our model to allow customers to have private information about the private-value component; with a few exceptions that we discuss in [Section 3.2](#), our conclusions remain unchanged as long as the amount of private information is small.) We assume that customers know the current state Z_t when they arrive to the market; that is, customers observe the event of news arrival but lack the expertise to interpret it. In light of this assumption, we can think of the parameter q as measuring the informational advantage of dealers over customers.

¹¹The signal we use is known as “truth-or-nothing,” and is commonly used in the literature on verifiable information following [Dye \(1985\)](#). In the context of financial markets it was used, for example, by [Baldauf and Mollner \(2020\)](#). It is also used in models of search with adverse selection, such as [Wolinsky \(1990\)](#) and [Blouin and Serrano \(2001\)](#).

¹²Given that the asset value is distributed symmetrically around zero, our analysis of the customer-buyer case can be translated to the customer-seller case by switching the roles of the high and the low type and changing the sign of prices.

2.2 Market protocols

We introduce four main market protocols that we will consider in our analysis.

The opaque market. In the opaque market, customers receive no public information about past or current prices, and must search to obtain a quote. Each customer can contact the first dealer and request a quote at no cost;¹³ any subsequent request for quote, from either dealer, incurs a search cost $s > 0$. Dealers are not committed to their quotes—they can provide a new quote if the same customer visits them again.¹⁴ We also assume that search is directed, that is, the customer can decide whether to visit dealer A or dealer B first. (To reduce spurious equilibrium multiplicity, we assume that, when the two dealers are indistinguishable from the perspective of the customer other than through their identity, the customer always visits her “default dealer” first, and the default dealer is equally likely to be dealer A or dealer B .) Upon being visited by a customer, the dealer quotes a price; then, the customer decides whether to accept (buy a unit of the asset at the quoted price) or reject. Conditional on rejecting, the buyer decides whether to search or exit the market without trading. A customer’s search history is their private information.

To reduce the number of market protocols to consider, we will only look at two limiting cases of search costs: (1) **low search costs**, in which s is low enough that customers decide to search whenever they would strictly prefer to search absent search costs, and (2) **high search costs**, in which s is high enough that customers never search in equilibrium. (Our proofs construct explicit bounds on s guaranteeing these equilibrium properties.)

The high-cost regime corresponds to a maximally opaque market in which each dealer is effectively a monopolist for its segment of the market. The low-cost regime models a market in which it is relatively easy to “shop around” for the best quote, but the sequentiality of search creates opacity on two levels; first, customers do not have access to both quotes at the same time, and second, dealers are not aware of the terms of trade offered by the competitor. The two search-cost regimes could be interpreted as capturing exogenous characteristics of the market but they could also be a consequence of regulation. For example, banning voice trading and enforcing an electronic request-for-quote protocol could be modeled as moving from the high-cost regime to the low-cost regime.

The post-trade transparent market. In the post-trade transparent market, the trading protocol is the same as in the opaque market but, at the end of each period t , a public signal is sent to all market participants specifying whether a transaction took place

¹³It is then a dominant strategy to request the first quote, and we assume that the customer always does it. If the first contact is costly, there always exist (uninteresting) equilibria with no trade.

¹⁴Lack of recall is a realistic feature of OTC search markets—see [Zhu \(2012\)](#). In our model, we show that in equilibrium no dealer is contacted more than once, and therefore our equilibrium characterization remains valid under search with recall.

at t , and if a transaction took place, then who was trading and at what price.¹⁵

Our modeling of the post-trade transparent market is inspired by regulation such as the introduction of TRACE in the US corporate bond market (see, for example, [Asquith et al., 2019](#)). The main consequence of such transparency is that traders can condition their pricing and trading strategies on past transaction data. In particular, investors can get a better sense of prices at which the asset has been trading recently, and dealers can refine the estimates of the their competitors trading positions and private information. Such regulation is often motivated as an attempt to “level the playing field” for smaller investors by giving them access to information that was only available to dealers in opaque markets.

The pre-trade transparent market. In the pre-trade transparent market, dealers compete directly by participating in an auction initiated when a buyer submits a request for quote to a trading platform. Formally, dealers A and B submit offers simultaneously in a first-price auction. The buyer can decide whether or not to trade, and with whom, after observing both offers. We assume that auction quotes become visible to all auction participants after the auction concludes.¹⁶ However, period- t customers observe nothing about the outcome of auctions in previous periods.

Our modeling of pre-trade transparency is inspired by regulation centralizing opaque OTC markets by requiring all trades to be executed on trading platforms. A prominent example is the introduction of Swap Execution Facilities that must be used in the US swap market following the Dodd-Frank act. The goal of the regulation is to increase competition between dealers by providing the quote requester with simultaneous access to all offers.

The fully transparent market. In the fully transparent market, there is both pre-trade and post-trade transparency. That is, trades happen via first-price auctions (as in the pre-trade transparent regime), and, at the end of each period, conditional on a transaction occurring, a public signal is sent to all market participants disclosing who traded and at what price (as in the post-trade transparent regime).

The fully transparent market has almost all features of a centralized trading protocol, typical for traditional stock exchanges. Traders have simultaneous access to all available quotes, and they observe past transaction data. However, even in this case, our framework retains one key feature of an OTC market: the privileged position of dealers in the trading network.

¹⁵At the end of Section 3, we comment on how the results change when past transaction data is summarized by a benchmark, thereby making the identities of dealers that traded unavailable to future customers.

¹⁶This assumption is satisfied in many but not all cases. [SIFMA \(2016\)](#) reviews the trading protocols used by various Swap Execution Facilities in the US, and reveals significant heterogeneity in the details of the auction designs. We comment on what happens in our model when dealers do not observe each others’ quotes at the end of Section 3.

2.3 Solution concept

An equilibrium is described by players' strategies and their beliefs at every information set. Specifically, under each of the protocols that we study, the strategy of dealer $i \in \{A, B\}$ consists of a history-contingent quote, and the period- t customer's strategy of her decision to accept one of the available quotes, search (if applicable) or exit the market. Our solution concept is symmetric stationary perfect Bayesian equilibrium with separating strategies for the dealers. By stationary equilibrium, we mean that, conditional on the hierarchy of beliefs about dealers' type (that also pins down all players' beliefs about the asset value), dealers' and customers' strategies can depend on calendar time t only through the (publicly observed) state Z_t . By a separating strategy for a dealer we mean that, on equilibrium path, the quote of dealer i at time t should uniquely pin down dealer i 's belief about the value of the asset at time t .

The restriction to equilibria with separating strategies for dealers is our strongest assumption. While $\Delta < 1$ rules out pooling equilibria (regardless of the regime), there may exist semi-pooling equilibria. The task of characterizing all such equilibria seems daunting. More importantly, we choose not to consider them for two reasons. First, prices in financial markets play two distinct roles: They directly affect trading decisions in the market and they serve as signals for actors in the real economy (Fama and Miller, 1972). By focusing on separating equilibria, we keep the second (unmodeled) welfare channel invariant across different regimes. If we found that some semi-pooling equilibria increase welfare in our model relative to separating equilibria, we would be implicitly introducing a trade-off since these equilibria would provide less information to the (unmodeled) outsiders relying on financial-market prices to make real-economy decisions. Second, the focus on separating equilibria seems natural in the context of transparency regulation, since the goal of the regulation, by definition, is to increase the amount of information revealed to the market; if market participants respond to the regulation by switching to an uninformative pricing strategy, the regulation does not achieve its goal (a point made elegantly by Kakhbod and Song, 2020).

Our environment is dynamic and features two-sided incomplete information. In the absence of equilibrium refinements on off-path beliefs, economically unnatural outcomes can be sustained by equilibria belonging to the class defined above. Equilibrium refinements that uniquely select a separating equilibrium have been widely used in the literature on signaling (Banks and Sobel, 1987, Cho and Kreps, 1987) and on adverse selection (Guerrieri et al., 2010, Guerrieri and Shimer, 2014). However, these refinements are not readily applicable in our environment which is dynamic and features multiple informed players. Instead of attempting to select a unique equilibrium, we impose relatively mild refinements that only rule out economically counterintuitive outcomes (we explain later how we deal with the resulting

equilibrium multiplicity). Specifically, we want to ensure that once the game transitions into “complete information” (in that all players are fully and symmetrically informed about the value of the asset), the continuation play (conditional on the state being persistent) is the efficient one that would arise in the analogous game of complete information.¹⁷ In the complete-information version of our protocols, dealers would quote a price equal to the consumer’s value in the pre-trade opaque market (the monopoly outcome) and a price equal to the value of the asset (the Bertrand outcome) in the pre-trade transparent market. The three refinements we describe next ensure that this is indeed how continuation equilibria look like after a fully revealing trade.

First, we require that the expected value of the asset is monotone in the observed quote. That is, for every player, the expected value of the asset after observing a quote b must be weakly increasing in b . This property must hold in equilibrium for on-path quotes b , so the assumption only has bite off-path.¹⁸ Second, we use a version of the “never dissuaded once convinced” restriction commonly used in bargaining models with incomplete information (see [Osborne and Rubinstein, 1990](#)). Specifically, we restrict the support of customers’ beliefs after observing an off-path event to be a weak subset of the support of her beliefs before observing said event. Finally, we require that, whenever a customer observes a pair of simultaneous quotes such that only one of them is consistent with on-path play, her posterior belief must be consistent with the quote that is on path (which only has bite when the on-path quote is fully revealing).¹⁹ In [Appendix OA.2](#), we show that if these refinements are not imposed, we can indeed construct some unreasonable equilibria; we also discuss the relationship to stronger refinements such as the D1 criterion by [Banks and Sobel \(1987\)](#).

2.4 Welfare

Since the gain from trade is always Δ , equilibrium welfare only depends on the expected probability of trade in every period. (We decide to exclude search costs from our measure of welfare since we think of search costs as a reduced-form way of modeling the market frictions, rather than as a physical cost incurred by the market participants.²⁰) Let $\underline{\Lambda}_r \in [0, 1]$

¹⁷This property is not guaranteed in a perfect Bayesian equilibrium because players are free to abandon their beliefs—even in case of common knowledge of a certain state—following any off-path event.

¹⁸This refinement is relatively standard in the literature; it is used in other dynamic signaling games such as [Daley and Green \(2012\)](#), [Gul and Pesendorfer \(2012\)](#), and [Halac and Kremer \(2020\)](#), and is similar in spirit to the refinements of [Cho and Kreps \(1987\)](#) and [Banks and Sobel \(1987\)](#) since the gain from reducing the quote (and thus inducing the customer to trade) decreases with dealers’ belief about the asset’s value.

¹⁹This refinement, which was introduced by [Bagwell and Ramey \(1991\)](#) under the name *unprejudiced beliefs* and whose theoretical properties are studied by [Vida and Honryo \(2021\)](#), is commonly used in signaling games with multiple senders. It is also close in spirit to the refinement proposed by [Milgrom and Mollner \(2021\)](#).

²⁰Given that we restrict attention to the case in which either the search cost is arbitrarily low or it is high enough to rule out search in equilibrium, including the search costs in the computation of welfare would not

and $\bar{\Lambda}_r \in [0, 1]$ denote the expected probability of trade, averaged and time-discounted over all periods, in regime r under the worst and best equilibrium, respectively, where $r \in \{\text{opaq}_{\text{low}}, \text{opaq}_{\text{high}}, \text{post}_{\text{low}}, \text{post}_{\text{high}}, \text{pre}, \text{full}\}$ indexes the six market regimes that we introduced in Section 2.2. We will be using the following notions of welfare dominance in face of equilibrium multiplicity.

Definition 1. Regime r weakly welfare-dominates regime r' if $\underline{\Lambda}_r \geq \underline{\Lambda}_{r'}$ and $\bar{\Lambda}_r \geq \bar{\Lambda}_{r'}$, with at least one inequality being strict. Regime r strongly welfare-dominates regime r' if $\underline{\Lambda}_r \geq \bar{\Lambda}_{r'}$ and $\bar{\Lambda}_r > \underline{\Lambda}_{r'}$.

Weak dominance of regime r over regime r' means that regime r leads to higher welfare both under the worst-case as well as the best-case equilibrium selection; however, it cannot be guaranteed that regime r will always lead to a better outcome than regime r' . Weak dominance is a compelling notion when all equilibria in a given regime look equally plausible to the regulator. In that case, a weakly-dominating regime provides the same (or better) worst-case guarantee, while allowing for a better (or not worse) outcome in the best-case scenario. Strong dominance is a more demanding but fully robust notion: If regime r strongly welfare-dominates regime r' , then *any* equilibrium under regime r outperforms even the best equilibrium under regime r' . Thus, a regulator has no justification for choosing r' over r .²¹

3 Welfare comparisons of trading regimes

In this section, we describe the structure of equilibria in each of the six regimes, and compare the regimes in terms of the expected welfare they generate. Throughout, we assume that

$$\delta\gamma \leq \bar{\delta}_{\text{exist}} \equiv \frac{(2-q)\Delta^2 - q}{(1-q)(1+3\Delta)}, \quad (3.1)$$

which, as we will show, guarantees equilibrium existence in all regimes (the condition for existence is more permissive for some of the regimes). Intuitively, when either the persistence γ in the asset value or dealers' patience δ is too high, a separating equilibrium may fail to exist since dealers may be able to benefit by deviating at t and manipulating the market's perception of the asset value at $t+1$. Recall that q —the probability of receiving an informative signal—measures the amount of adverse selection in the market. Thus, the condition requires that the amount of adverse selection should not be too large relative to Δ —the

affect our welfare rankings. This is because, under this restriction, the total search cost that is incurred in equilibrium is either zero (under the high search cost regime) or negligible (under low search costs).

²¹The reason why we only focus on the worst and the best equilibrium is that any level of welfare between these two extremes can also be achieved as long as a public randomization device is available.

private-value “buffer” between dealers and customers. Overall, our modeling assumptions combined with the above restriction on parameters imply that our analysis is relevant for OTC markets in which the asset value changes relatively frequently (so that any informational advantage is short-lived), the amount of private information about the asset value is not excessive, and customers trade primarily for liquidity reasons.

3.1 The opaque market

Our first result characterizes equilibria in the opaque regime with high search costs. Because high search cost effectively makes each dealer a monopolist to the visiting customers, and post-trade opacity kills any information spillovers across periods, it is instructive to first consider a static monopoly benchmark in which there is only one dealer, the trade protocol is request for quote (the dealer quotes a price and the customer accepts or rejects), and trade happens only once (there is no continuation game).

In a separating equilibrium of the static game with a single dealer, a quote b_σ of type $\sigma \in \{l, u, h\}$ reveals the value of the asset to the customer; at the same time, the dealer with type σ must find it optimal to quote b_σ . Since in equilibrium we must have $b_l < b_u < b_h$, it follows immediately that types u and h cannot trade with probability one, as otherwise the low-type dealer would always have a profitable deviation. If types u and h trade at all, it must therefore be the case $b_u = \Delta$ and $b_h = 1 + \Delta$ since only these quotes make the consumer indifferent between accepting and rejecting (making it possible to implement an interior probability of trade by having the customer mix in equilibrium). Finally, we argue that $b_l = -1 + \Delta$ and the quote must be accepted with probability one. Since the customer’s expected value for the good is at least $-1 + \Delta$ (regardless of beliefs), the low type can guarantee a profit of $\Delta - \epsilon$, for any $\epsilon > 0$, by quoting $-1 + \Delta - \epsilon$ (the customer strictly prefers to accept this quote). On the other hand, a quote $b_l > -1 + \Delta$ would be rejected by the customer. It follows that the low type must make a profit of Δ in equilibrium with a quote $b_l \leq -1 + \Delta$, which is only possible when $b_l = -1 + \Delta$.

In equilibrium of the static game, the following dealer incentive-compatibility constraints must hold, ruling out profitable deviations of mimicking the behavior of a different type:

$$\begin{aligned} \Delta &\geq \lambda_u(\Delta - (-1)), & (\text{IC}_{l \rightarrow u}) \\ \Delta &\geq \lambda_h((1 + \Delta) - (-1)), & (\text{IC}_{l \rightarrow h}) \\ \lambda_u \Delta &\geq \lambda_h((1 + \Delta) - 0), & (\text{IC}_{u \rightarrow h}) \end{aligned}$$

where λ_u and λ_h are the customer’s equilibrium probabilities of accepting the quotes b_u and b_h , respectively. In the proof of Proposition 1 in Appendix B, we complete this characteriza-

tion by showing that any (λ_u, λ_h) satisfying these conditions correspond to an equilibrium, and that all equilibria are payoff-equivalent to some equilibrium of the form described above.

We are ready to characterize equilibria of the opaque high-search-cost market: Under our solution concept, they reduce to repetitions of the static monopoly trading pattern in every dealer-customer interaction.

Proposition 1 (Opaque high-search-cost market). *In all equilibria, in each period t , equilibrium outcomes in the opaque high-search-cost regime are payoff-equivalent to the static monopoly trading pattern: The low type quotes $-1 + \Delta$ which is accepted with probability one, the uninformed type quotes Δ which is accepted with probability $\lambda_u \in [0, \bar{\lambda}_u]$ for some $\bar{\lambda}_u < 1$, and the high type quotes $1 + \Delta$ which is accepted with some probability $\lambda_h \in [0, \bar{\lambda}_h]$ for some $\bar{\lambda}_h < \bar{\lambda}_u$.*

In the least efficient equilibrium, the uninformed and the high-type dealers do not trade. In the most efficient equilibrium, in each period, the uninformed type trades with probability $\Delta/(1 + \Delta)$, and the high type trades with probability $(\Delta/(1 + \Delta))^2$.

The equilibria described in Proposition 1 have a simple structure: The customer does not search, so each dealer is effectively a monopolist. Because there is no informational link between periods, trading is unaffected by what happened in past periods. In the least efficient equilibrium, the low and the high type do not trade at all, so only a fraction $q/2$ of the first-best welfare is realized. Even in the most efficient equilibrium, $\lambda_u = \Delta/(\Delta + 1)$ and $\lambda_h = (\Delta/(\Delta + 1))^2$, and hence a significant fraction of gains from trade are lost.

The relatively low efficiency of trading in the opaque high-search-cost market is a consequence of the information asymmetry between dealers and customers. In our model, this inefficiency is manifested through the randomized acceptance decision by the customer. Customers' mixing is a robust feature of all the equilibria we construct in the remainder of this section; it is important to emphasize that it is not a technicality—rather, it is shaped by fundamental economic forces. Due to adverse selection and signaling, dealers post unattractive prices that make the customer just indifferent between trading or not. Then, in equilibrium, the buyer accepts with probabilities that make dealers' pricing decisions optimal.²² The lower and upper bounds on customers' trading probabilities are pinned down by the binding

²²In Appendix OA.1, we consider an extension of our model in which period- t customer's private value for trading is $\Delta + \epsilon_t$, where ϵ_t is a mean-zero continuous random variable whose realization is privately observed by the customer. Then, acceptance decisions are deterministic, pinned down by a cutoff in the space of realizations of ϵ_t . Intuitively, our simple model corresponds to the limiting case in which the variance of ϵ_t converges to zero. Even when the uncertainty over ϵ_t disappears, dealers do not adopt pricing strategies that always lead to a transaction, due to adverse selection. In contrast, if the common value component were common knowledge among all traders, the probability of trade would be below 1 for any fixed non-degenerate ϵ_t but would converge to 1 with vanishing uncertainty over ϵ_t .

incentive-compatibility constrains for the dealers, reflecting the informational advantage of the dealers over the customers.

A natural conjecture is that the low welfare in the opaque market is driven by lack of competition between the dealers in the face of high search costs. However, as we show next, even when customers can search effectively for free, welfare remains low at least in some (and sometimes in all) equilibria. We call an equilibrium a *no-search equilibrium* if the customer visits exactly one dealer on equilibrium path in every period t .

Proposition 2 (Opaque low-search-cost market). *Regardless of the search cost, there always exists a no-search equilibrium in which dealers with a high or uninformed type never trade. This is the least efficient equilibrium outcome.*

In the opaque low-search-cost regime, if $q \leq 2\Delta^2/(1 + \Delta)$, there exist equilibria with search in which the customer searches following a quote from a high-type dealer when $Z_t = 0$ (and does not search in any other case), and an equilibrium of this form is the most efficient equilibrium. If $q > 2\Delta^2/(1 + \Delta)$, the unique equilibrium is the no-search equilibrium in which dealers with a high or uninformed type never trade.

Proposition 2 predicts that the highly inefficient no-search equilibrium can be sustained even when search costs are arbitrarily low. To understand why, note that, in equilibrium, the customer always makes zero profits when trading against an informed dealer given the sequentiality of search. Thus, the only case in which search can be beneficial for the customer is when the first quote she receives is high, and the second quote is coming from an uninformed dealer (who then underprices the asset relative to its true value). In fact, the customer will always want to search following the high quote provided that her search costs are low enough and the uninformed type provides a “serious” quote.

However, in the worst equilibrium, the uninformed dealer does not provide a serious quote. Suppose that the customer never buys from the uninformed dealer if this is the first dealer she visits—such an outcome can be supported in equilibrium because the quotes always make the customer indifferent between accepting or not. Then, if there is search in equilibrium, the uninformed dealer faces an extreme adverse selection problem: She sells *only* when the customer visits a high dealer before visiting her, in which case the value of the asset must be high. Thus, it is strictly suboptimal for the uninformed dealer to offer any quote below $1 + \Delta - s$. Instead, she offers a high quote $1 + \Delta - \epsilon$, for some sufficiently small ϵ , which is always rejected. It then becomes optimal for the customer to never search; in equilibrium, the customer must also reject the quote from the high type, as otherwise the uninformed dealer would be able to profitably deviate by quoting a high price.

The more efficient search equilibrium exists if adverse selection (as measured by q) is not too severe for the uninformed dealer. Suppose that the uninformed type quotes Δ . She

can then make positive profits if she is first to be visited, but she makes losses when she is second to be visited and the first dealer visited by the customer has a high type. The condition $q \leq 2\Delta^2/(1 + \Delta)$ ensures that the uninformed type at least breaks even with this strategy. The search equilibrium is clearly more efficient than the no-search equilibrium since the uninformed type trades with positive probability.

Still, the value of search is limited. We show that if there is search by the customer in period t , then there cannot be search in period $t + 1$ if $Z_{t+1} = 1$ (the state is persistent). This is due to a subtle learning effect. An uninformed dealer realizes that the customer is more likely to buy from her when the asset value is high (because the customer buys for sure if she received a high quote before visiting an uninformed dealer). Thus, an uninformed dealer becomes more optimistic about the value of the asset when she is visited and sells. If the state is persistent, the customer who arrives in the following period will search not only after seeing a high quote, but also after seeing the quote of the “optimistic” uninformed type, in the hope that she will then encounter the more pessimistic uninformed type. However, this would mean that the “optimistic” uninformed type is subject to the extreme form of adverse selection: She sells only when she is visited after a visit to a high-type dealer, that is, only when she underprices the asset relative to its true value. This cannot be an equilibrium. As a result, following search, the subsequent period features no search if the state is persistent, with uninformed dealers offering high quotes that are always rejected.

3.1.1 Is the opaque market more efficient when search cost goes down?

We now turn to the welfare comparison associated with lowering the search costs.

Proposition 3 (Value of search in opaque markets). *There exist cutoffs \underline{q}, \bar{q} , with $0 < \underline{q} < \bar{q} < 1$, such that the low-search-cost opaque market weakly welfare-dominates the high-search-cost opaque market if $q \in (\underline{q}, \bar{q}]$, and the high-search-cost opaque market weakly (strongly) welfare-dominates the low-search-cost one if $q < \underline{q}$ ($q > \bar{q}$).*

Proposition 3 predicts that the effects of lowering search costs in an opaque market are ambiguous. The least efficient equilibrium leads to the same welfare in both cases; the comparison is thus driven by what happens in the most efficient equilibria.

When q is high enough ($q > \bar{q}$), our results in Propositions 1 and 2 imply that there is never search in equilibrium, no matter how low the search cost is. In particular, the low-search-cost setting has a unique equilibrium outcome in which only the low type trades, which coincides with the worst-case equilibrium outcome in the high-search-cost environment. However, under high search costs, it is possible to sustain other more efficient no-search equilibria in which the uninformed and high types trade with strictly positive probability. Thus, a higher search cost leads to strong welfare dominance.

For intuition, recall that the only way to sustain no search as an equilibrium outcome when s is low is if the uninformed type (and as a result the high type) provides a “non-serious” quote that leads the customer to exit the market without trading. Otherwise, the uninformed type would not be able to break even due to adverse selection. In contrast, when the search cost is high, dealers act as monopolists upon being contacted by a customer and the uninformed dealer can sell with positive probability. The high search cost protects the uninformed dealer from the search-induced adverse selection.

When $q \leq \bar{q}$, the equilibrium outcome in the low-search-cost setting is no longer unique and the most efficient equilibrium features search on equilibrium path. We show that the difference in welfare levels sustainable in the best equilibria of the two regimes is monotone in q and crosses zero at some interior value of $q \in (0, \bar{q})$. Intuitively, the low-search-cost regime is more efficient at $Z_t = 0$ (due to search) but less efficient at $Z_t = 1$ (since the uninformed and high types do not trade in this case); as q gets small, probability of trade at $Z_t = 1$ converges to zero, outweighing the welfare gain conditional on $Z_t = 0$.

Overall, lowering the costs of search in an otherwise opaque market with adverse selection can have subtle—and potentially negative—equilibrium effects. On the one hand, search allows the customer to gather information, which generates an adverse selection problem for the uninformed dealers. As a result, when q is sufficiently high, search-induced adverse selection forces uninformed dealers to provide unattractive quotes that are always rejected. On the other hand, when the adverse selection problem is not too severe (so that uninformed types can trade at a profit even if the customer searches), the low-search-cost setting yields a higher probability of trade than the high-search-cost setting in the event that dealers are informed. As a result, the former weakly welfare-dominates the latter when q takes on intermediate values. For lower values of q , the welfare ranking is reversed and coincides with the ranking obtained when q is high.

3.2 The post-trade transparent market

We next turn to the characterization of equilibria in the post-trade transparent regime, with both high and low search costs. The difference to the opaque market is that a transaction taking place in period t becomes publicly announced to all future market participants. When we refer to the “static monopoly trading pattern,” we mean the within-period equilibrium structure defined in Proposition 1.

Proposition 4 (Post-trade transparent high-search-cost market). *If $Z_t = 0$, the static monopoly trading pattern emerges. If $Z_t = 1$:*

- *If a low-type or high-type dealer transacted in period $t - 1$, the equilibrium is fully*

efficient in period t ;

- *If an uninformed dealer transacted in period $t - 1$, then both that dealer and the low-type of the other dealer trade with probability one conditional on being contacted by a customer in period t ; the remaining types trade with probability strictly less than one;*
- *If there was no transaction in period $t - 1$, then the static monopoly trading pattern emerges in period t .*

Proposition 4 is intuitive. If the state is persistent at t , the customer and the dealers acquire some information from the public disclosure of the previous trade, which then affects equilibrium strategies.

If an informed dealer transacted in period $t - 1$, then equilibrium in period t becomes fully efficient. This case captures the intuitive appeal of post-trade transparency: The informational content of prices becomes publicly revealed, and market participants become symmetrically informed. Thus, trade also becomes fully efficient, even when there is no direct competition between the dealers.²³

The benefits of post-trade transparency are more limited in the other two cases. When the uninformed dealer trades in period $t - 1$, the public announcement conveys no information about the value of the asset but it creates common knowledge that the dealer who transacted is uninformed. Hence, upon being contacted by a customer in period- t , that dealer has no need to signal what her type is, which allows her to trade at the surplus-extracting price without any inefficiency. On the other hand, the other dealer's type is still unknown to the customer in period t , so trade against that dealer remains inefficient. Since the customer makes no profits either way, in equilibrium she can direct search towards either dealer; this creates additional equilibrium multiplicity and has consequences for welfare (which we will describe in the next subsection). Finally, post-trade transparency offers no benefits in case no transaction occurs in period $t - 1$.²⁴

Next, we describe equilibria in the post-trade transparent market with low search costs.

Proposition 5 (Post-trade transparent low-search-cost market). *Regardless of the search cost, there always exists a no-search equilibrium in which dealers with a high or uninformed type never trade, except when there is trade at the low price in period $t - 1$ and $Z_t = 1$ (in which case trade happens with probability one in period t). This is the least efficient equilibrium outcome.*

²³Even though trade is efficient in this case, dealers extract all the surplus. Thus, “leveling the playing field” through post-trade transparency alone is not guaranteed to benefit the customers.

²⁴Recall that—even though we focus on the case that the customer is a buyer in describing the equilibrium—the customer is a seller with equal probability. This is why the event of no trade is equally likely to occur under both high and low asset value.

In the post-trade transparent low-search-cost regime, if $q \leq \bar{q}_{\text{post low}}$, where $\bar{q}_{\text{post low}} \equiv 2\Delta^2(1 + \delta\gamma/2)/((1 + (1 - \delta\gamma/2)\Delta)(1 + \delta\gamma\Delta/2))$, there exist equilibria with search in which the customer searches following a quote from a high-type dealer when $Z_t = 0$ (and does not search in any other case), and an equilibrium of this form is the most efficient equilibrium. If $q > \bar{q}_{\text{post low}}$, the unique equilibrium is the no-search equilibrium.

The first part of Proposition 5 shows that post-trade transparency does not eliminate the inefficiency caused by adverse selection even when search costs are very small. The reason is intuitive. Post-trade transparency reduces information asymmetries only when there is a transaction to be disclosed. However, in the worst equilibrium, the uninformed and the high-type do not trade at all. Some improvement is achieved though: Because the low type always sells, the period- t outcome (conditional on the state being persistent) is fully efficient following a transaction at a low price in period $t - 1$.

When the adverse selection problem is not too severe, there exist equilibria with search, similar to those discussed in Proposition 2. But once again, the role of search is limited by the adverse selection problem. In particular, we show that, if there is search at $Z_{t-1} = 0$, there cannot be any search at $Z_t = 1$ following a transaction in period $t - 1$. This prediction is similar to the one made for the opaque market (see Proposition 2) but the reasons are different. Conditional on a high or low price in period $t - 1$, there is no search at $Z_t = 1$ because the market becomes fully efficient. Conditional on an uninformed type trading in period $t - 1$, it becomes common knowledge that the dealer who traded is uninformed. In an equilibrium with search, the period- t customer would then prefer to visit the other dealer first, hoping that that dealer's type is high, and she can visit the uninformed dealer next and make a profit. This cannot happen in equilibrium, though, because the uninformed dealer would only trade when the value of the asset was high, making a loss. Instead, in equilibrium, the uninformed dealer who traded in period $t - 1$ provides an unattractive quote at $Z_t = 1$, and is never visited by the customer (hence, there is no search).

3.2.1 Does introducing post-trade transparency increase market efficiency?

Propositions 4 and 5 predict that equilibria with post-trade transparency have a similar structure to the ones in the opaque market; however, they lead to fully efficient trade at $Z_t = 1$ in at least some contingencies.

Proposition 6 (Post-trade dominates opaque). *The post-trade transparent market weakly welfare-dominates the opaque market (under both high and low search costs). Under low search costs, there exists $\bar{q} \in (0, 1)$ such that, if $q > \bar{q}$, the post-trade transparent market strongly welfare-dominates the opaque market.*

Proposition 6 may seem intuitive given the preceding discussion since the fully efficient outcome following an informative trade raises welfare in both the best and worst equilibrium. However, the result is not *a priori* obvious because the introduction of post-trade transparency creates additional incentives for dealers to mimic a higher type when $Z_t = 0$. The reason is that, by quoting a higher price, dealers can now manipulate the beliefs of customers in period $t + 1$ (conditional on their quote being executed). This in turn leads to tighter incentive-compatibility constraints for the dealers, which may hinder welfare.²⁵

The reason why this negative aspect of post-trade transparency is not manifested in Proposition 6 is that post-trade transparency has an additional advantage: It allows future customers to condition their search strategy on past trading outcomes. In particular, the customer in period $t + 1$ may condition her choice of which dealer to contact first on who traded in period t . In the best equilibrium, the customer breaks her indifference between dealers at $Z_{t+1} = 1$ in a way that relaxes dealers' incentive constraints in period t . In other words, the benefits of manipulating future customers' beliefs upwards can be eliminated if equilibrium specifies that future customers (in case of indifference) are more likely to visit dealers that traded at lower prices in the recent past. This effect also explains why post-trade transparency strongly welfare-dominates the opaque market in the case of low search cost and intermediate values of q : Post-trade transparency allows to sustain the more efficient search equilibrium for higher levels of adverse selection compared with the opaque market.²⁶

The fact that the positive effect described above dominates the negative effect of belief manipulation is sensitive to our assumption that the customer has no private information. In Appendix OA.1, we study an extension of the model in which customers' values are perturbed in a way that breaks their indifference. Intuitively, with a little bit of private information about their values for trading, customers strictly prefer to visit dealers with whom they have a higher ex-ante probability of trading. This modification turns out to have no effect on other regimes but it limits the extent to which revealed trades can guide search decisions of future customers. As a result, the post-trade transparent market no longer weakly welfare-dominates the opaque market in all the cases—we show that there are values of q for which the two regimes cannot be ranked; post-trade transparency increases welfare in the worst equilibrium, but lowers welfare in the best equilibrium (due to the belief-manipulation effect).

Even in our baseline model, the post-trade transparent market typically does not strongly welfare-dominate the opaque market. This means that the best equilibrium in the opaque market may be better than the worst equilibrium in the post-trade transparent market.

²⁵A similar effect, albeit in a different context and setting, appears in Dworzak (2020) and Antill and Duffie (2020).

²⁶For high values of q , post-trade transparency strongly welfare-dominates the opaque market with low search costs because the latter has a unique highly inefficient equilibrium.

Thus, post-trade transparency regulation (such as TRACE) may lower welfare if market participants coordinate on an inefficient equilibrium. The only case in which the introduction of post-trade transparency leads to an unambiguous improvement in market efficiency is when the original opaque market has low search costs and high adverse selection.

Finally, we note that the effects of lowering the search costs in the post-trade transparent market are ambiguous, just like in the opaque market.

Proposition 7 (Value of search in a post-trade transparent market). *There are cutoffs \underline{q}, \bar{q} , with $0 < \underline{q} < \bar{q} < 1$, such that the high-search-cost post-trade transparent market weakly welfare-dominates the low-search-cost post-trade transparent market if either $q < \underline{q}$ or $q \geq \bar{q}$.*

3.3 The pre-trade transparent market

In this section, we consider the effects of changing the trading protocol from search to direct competition between dealers on a trading platform.

Proposition 8 (Pre-trade transparent market). *In the pre-trade transparent market, in all equilibria, when $Z_t = 0$, the high type quotes 1, the uninformed type quotes a deterministic price $b_u \leq \Delta$, and the low type quotes a random price distributed on $[\underline{b}, -1 + \Delta]$, with $\underline{b} > -1$. Given a pair of on-path quotes (q_A, q_B) , the customer accepts the lower quote with probability one; if $q_A = q_B$, the customer breaks the tie between dealers uniformly at random but refuses to trade with probability $1 - \lambda_u$ when $q_A = q_B = b_u$.*

When $Z_t = 1$, trade in period t is fully efficient.

There exists a fully efficient equilibrium (with $\lambda_u = 1$) if and only if $\delta\gamma \leq \bar{\delta}_{\text{pre}} < \bar{\delta}_{\text{exist}}$, where $\bar{\delta}_{\text{pre}} > 0$ if and only if $q < (2\Delta - 1)/(2\Delta)$.

Proposition 8 describes the structure of equilibria with pre-trade transparency. The first notable effect is that the direct competition between dealers shifts the gains from trade towards the customers. In particular, high-type dealers price the asset at its expected value. Moreover, trade is fully efficient whenever at least one of the dealers is informed. The only case in which trade may not take place is when both dealers are uninformed and $Z_t = 0$. We show in the proof that there are two possible trading outcomes that may arise in this event.

First, there always exists an equilibrium where the uninformed type quotes $b_u = \Delta$. In this equilibrium, when the buyer faces prices (Δ, Δ) , she is indifferent between buying or not, and she sometimes rejects both offers. The probability of accepting is bounded above because otherwise the low-type dealer would find it profitable to deviate and quote Δ . However, unlike in previous cases, the probability of accepting is also bounded from below. To understand why, note that the uninformed dealer must at least break even, and she is

making strict losses whenever she wins the auction against the high type (this is the classical “winner’s curse” effect). Thus, she must sell the asset sufficiently often in case she is tied against another uninformed dealer.

Second, there may exist an additional equilibrium in which $b_u < \Delta$ and trade happens with probability one under all equilibrium contingencies. This fully efficient outcome is fragile: It requires low adverse selection (low q), high private-value buffer (high Δ ; in particular above $1/2$), and low dealer patience and asset value persistence (low δ and γ ; in particular lower than required for equilibrium existence).

In either type of equilibrium, if the state is persistent at t , then period t features efficient trade, even though customers have no information about the value of the asset. The reason is that the dealers become symmetrically informed following the period- $(t-1)$ auction (since we assumed that quotes in the auction are observed). This, combined with direct competition, is sufficient to ensure efficiency—the dealers engage in Bertrand competition.

3.3.1 Which is better: post- or pre-trade transparency?

Perhaps unsurprisingly given the preceding discussion, we can show that pre-trade transparency outperforms post-trade transparency.

Proposition 9 (Pre-trade dominates post-trade). *Pre-trade transparency weakly welfare-dominates post-trade transparency (with either high or low search cost). When q is sufficiently high, pre-trade transparency strongly welfare-dominates post-trade transparency (with either high or low search cost).*

The proof shows that pre-trade transparency is strictly better than post-trade transparency both in the worst- and best-equilibrium, regardless of the parameters. But if adverse selection is severe, an even stronger conclusion obtains: The worst equilibrium in the pre-trade transparent regime is strictly better than the best equilibrium in the post-trade transparent regime. Intuitively, pre-trade transparency increases market efficiency conditional on the high type realizing, which is more likely when q is high. The high type cannot trade with probability one in a pre-trade opaque market because this would make lower types imitate the offers made by the high type. In the auction market, mimicking the quote of a high type is not profitable; a high quote can win only against another high quote, in which case the value of the asset is high, and the profit of the dealer is zero.

3.4 The fully transparent market

Finally, we turn attention to the fully transparent market in which trades in the auction become publicly revealed.

Proposition 10 (Fully transparent market). *Equilibria in the fully transparent market have the same structure as in the case of a pre-trade transparent market. However, the fully efficient equilibrium exists if and only if $\delta\gamma \leq \bar{\delta}_{\text{full}} < \bar{\delta}_{\text{pre}}$, and, in equilibria in which the uninformed dealers quote $b_u = \Delta$, the highest λ_u that can be sustained is strictly lower than in the pre-trade transparent market.*

The main conclusion of Proposition 10 is that adding post-market transparency on top of pre-trade transparency actually hurts market efficiency. There are two pieces of intuition for this result. First, post-trade transparency does not help in any contingency, since—as we emphasized previously—it is not actually important to make the customers more informed when dealers compete directly and are symmetrically informed. Second, post-trade transparency actually makes things worse because it tightens the incentive-compatibility constraints for the dealers who can now benefit from manipulating future customers’ beliefs.

To see why, consider a low-type dealer deviating to a quote of Δ , and assume the other dealer also quoted Δ . If there is no post-trade transparency, the informational effect of this deviation is that the other dealer will not learn the true value of the asset. But with post-trade transparency, an additional effect is that the next-period customer—having observed the winning quote of Δ —will believe that either both dealers are uninformed or one of them has a high type. Thus, the customer will be willing to accept any price below Δ . This allows the low-type dealer to make additional profits by undercutting the uninformed dealer in the following period (this is profitable, since the dealer knows that the actual value of the asset is -1). For contrast, without post-trade transparency, undercutting the uninformed type would not be beneficial for the low type because her quote would be rejected by an uninformed customer in equilibrium. Since this dynamic deviation becomes more attractive for the low-type dealer in the fully transparent regime, the probability of trade following the pair of quotes (Δ, Δ) must be lower to make the deviation unprofitable. At the same time, the uninformed dealer cannot similarly benefit from manipulating customers’ beliefs upwards because undercutting the high type is never profitable. Therefore, the uninformed type must trade with lower probability in the most efficient equilibrium (and the same probability in the least efficient equilibrium) when post-trade transparency is added.

3.4.1 Is full transparency beneficial?

We summarize the above discussion in the following result.

Proposition 11 (Pre-trade transparency dominates full transparency). *If $\delta\gamma \leq \bar{\delta}_{\text{full}}$, pre-trade transparency and full transparency have the same set of equilibrium levels of welfare. Otherwise, pre-trade transparency weakly welfare-dominates full transparency.*

The proof shows that welfare in the worst equilibrium is the same across the two regimes. When $\delta\gamma$ is below $\bar{\delta}_{\text{full}}$, both regimes admit a fully efficient equilibrium, and hence welfare in their best equilibrium is also the same. Since $\bar{\delta}_{\text{full}} < \bar{\delta}_{\text{pre}}$, and $\bar{\delta}_{\text{pre}} > 0$ only under relatively restrictive conditions, this case is very special. In all other cases, full transparency leads to lower welfare than pre-trade transparency in the best equilibrium. Thus, introduction of post-trade transparency to a pre-trade transparent market can only hurt.

3.5 Other variants of the trading regime

In the preceding sections, we proposed a simple taxonomy of trading regimes that classifies them based on the types of transparency and search frictions that they entail. The resulting six protocols allow us to identify, in a stylized way, the main channels through which transparency can affect equilibrium market outcomes. In this section, we briefly discuss two other salient trading regimes and argue that their analysis would lead to qualitatively similar conclusions.

Benchmarks. A transparency policy commonly observed in practice is the publication of a benchmark, which reveals information about some aggregate market outcome—typically, an average price in past transactions (Duffie and Stein, 2015). Compared to full post-trade transparency, a benchmark has two distinctive features. First, it makes dealers anonymous, since it suppresses the information about the identities of traders. Second, by averaging over many transactions, a benchmark reduces the impact that any given dealer has on future market beliefs. Because our model has at most one transaction in every period, it is difficult to capture the second feature without modifying the framework. But we can analyze a protocol that captures the first feature of benchmarks by assuming that past transactions are publicized without disclosing the identity of the traders. We will argue that the resulting equilibria are similar, and generally less efficient than in the post-trade transparent regime.

Regardless of the search cost, the least efficient equilibrium under a benchmark is the same as under post-trade transparency since our construction of the worst equilibrium did not rely on customers observing the identities of dealers who traded in previous periods. However, benchmarks have a drawback that lowers the best-case level of welfare relative to the post-trade transparent regime. Because trades are disclosed anonymously under the benchmark, the two dealers are symmetric from the perspective of the customer who arrives to the market in the persistent state—the customer is not able to choose which dealer to contact first based on past trading outcomes. Thus, unlike in the post-trade transparent regime, customer’s directed search decisions cannot be used to discipline dealers and disincentivize belief manipulation. For example, in the low-search-cost environment, if a low-type dealer quotes a low price in the post-trade transparent regime, then in the best equilibrium

she is “rewarded” by being visited by a customer in the next period with probability one (conditional on the state being persistent). But, under a benchmark, the customer cannot distinguish the two dealers, and hence the low-type is visited with conditional probability half in the following period after playing her equilibrium quote. This makes deviating to the uninformed type’s price more attractive; to prevent it, the uninformed dealer must trade with lower probability, which lowers overall welfare.

Summarizing, a benchmark—modeled in the particular way described above—would be weakly welfare-dominated by post-trade transparency. Because a benchmark would induce efficient outcomes following informative trades, it would weakly welfare-dominate the opaque market. Thus, similarly to [Duffie et al. \(2017\)](#), we find reasons to introduce benchmarks in fully opaque markets; however, benchmarks do not help close the welfare gap between a post-trade transparent market and a pre-trade transparent market.

Opaque auctions in the pre-trade transparent market. We assumed that dealers observe each others’ quotes in the pre-trade transparent regime; in fact, this feature played an important role in our analysis. In practice, whether or not rejected quotes in the auction are observable depends on the trading platform ([SIFMA, 2016](#)).

Allowing for the possibility that auction quotes are not observable to dealers (an “opaque auction”) raises several difficulties that make the analysis considerably less tractable. The reason is that, when rejected quotes are unobservable, dealers will be able to partially infer their opponent’s type at the end of each period by observing whether or not their quote won the auction. Moreover, dealer i ’s posterior beliefs about the opponent’s type will depend on i ’s quote and on the details of dealers’ quoting strategies at $Z_t = 0$. As a result, there is a continuum of different continuation games when $Z_t = 1$ that arise for each realization of dealers’ (privately observed) quotes in the previous period. In particular, because dealers will now continue to face asymmetric information when $Z_t = 1$, the efficient trading outcome is no longer the continuation equilibrium, suggesting that this modification cannot raise surplus compared to the baseline pre-trade transparent regime.

Additionally, an opaque auction is similar to a pre-trade opaque market with a zero search cost. (In fact, it would be equivalent if we allowed for recall in the search market.) Because an opaque-auction protocol can be seen as a hybrid between pre-trade transparency with observable quotes and pre-trade opacity with low search costs, we expect its welfare outcomes to lie somewhere in between as well.

4 Optimal trading mechanism

In this section, we use mechanism-design tools to construct a trading mechanism that maximizes welfare in our setting. Our approach is to impose relatively stringent requirements on feasible mechanisms—less permissive than the ones implied by the solution concept used in Section 3—and show that the optimal mechanism nevertheless outperforms all of the regimes analyzed so far.

By the revelation principle, we restrict attention to direct mechanisms in which agents report their exogenous private information to the mechanism, and the mechanism specifies the trading probabilities and transfers. We study ex-post implementation, requiring that agents' reports and participation decisions remain optimal even conditional on learning other agents' types. This restriction is a strengthening of the assumption that dealer strategies are fully revealing of their types. Additionally, ex-post implementation rules out the fully-efficient but economically unappealing mechanisms of Crémer and McLean (1985). Finally, to parallel the stationary structure of equilibria we constructed in Section 3, we assume that the mechanism can only condition on information that is relevant for the current asset value.²⁷

To set up the problem formally, we let $\theta_{i,t} \in \Theta \equiv \{-1, 0, 1\}$ denote dealer i 's type in period t with $Z_t = 0$. Note that we redefined the types so that they are equal to the corresponding expected value of the asset given the exogenous signal the dealer receives about v_t —this will be convenient when describing the payment rule in the optimal mechanism. Given that dealers do not receive any exogenous information in periods in which the state is persistent, we set $\theta_{i,t} = \theta_{i,t-1}$ when $Z_t = 1$. We let $v(\theta_i, \theta_{-i}) \equiv \mathbb{E}[v_t | \theta_{i,t} = \theta_i, \theta_{-i,t} = \theta_{-i}]$. Since the state Z_t was assumed common knowledge, we assume that it is also observable to the designer (as will be clear from our results, this assumption is not needed to implement the optimal mechanism). A fully general dynamic mechanism conditions outcomes at time t on the entire history of reports and states up to and including time t . We instead restrict the designer to use a *stationary mechanism* that conditions outcomes at time t only on period- t reports, the realization of Z_t , and reports in $t - 1$ if $Z_t = 1$. Because the only players who possess private information are dealers at times t when $Z_t = 0$, we can further restrict attention to mechanisms in which dealers report their types to the mechanism only when $Z_t = 0$. Overall, given that a symmetric mechanism is without loss of generality in our setting, we define a trading mechanism to be $(x_Z, p_Z, b_Z) : \Theta^2 \rightarrow [0, 1] \times \mathbb{R}^2$, for $Z \in \{0, 1\}$, where $x_Z(\theta_{i,t}, \theta_{-i,t})$ is dealer i 's probability of trade at time t when $Z_t = Z$, i reported $\theta_{i,t}$, and $-i$ reported $\theta_{-i,t}$.

²⁷This rules out mechanisms that can punish a dealer for misreporting their type by changing the future trading protocol—a feature that not only would be ruled out by the equilibrium concept from Section 3 but also seems unrealistic.

(in the period in which information arrived); $p_Z(\theta_{i,t}, \theta_{-i,t})$ is the corresponding monetary transfer that dealer i receives, and $b_Z(\theta_{i,t}, \theta_{-i,t})$ is the transfer paid by the customer (buyer). A mechanism is *feasible* if it satisfies:

- *Dealers' ex-post incentive compatibility:*

$$p_0(\theta_i, \theta_{-i}) - x_0(\theta_i, \theta_{-i})v(\theta_i, \theta_{-i}) + \delta\gamma(p_1(\theta_i, \theta_{-i}) - x_1(\theta_i, \theta_{-i})v(\theta_i, \theta_{-i})) \geq p_0(\theta'_i, \theta_{-i}) - x_0(\theta'_i, \theta_{-i})v(\theta_i, \theta_{-i}) + \delta\gamma(p_1(\theta'_i, \theta_{-i}) - x_1(\theta'_i, \theta_{-i})v(\theta_i, \theta_{-i})), \quad \forall \theta_i, \theta'_i, \theta_{-i} \in \Theta;$$

- *Dealers' ex-post per-period individual rationality:*

$$p_Z(\theta_i, \theta_{-i}) - x_Z(\theta_i, \theta_{-i})v(\theta_i, \theta_{-i}) \geq 0, \quad \forall \theta_i, \theta_{-i} \in \Theta, Z \in \{0, 1\};$$

- *Customers' ex-post individual rationality:*

$$(v(\theta_A, \theta_B) + \Delta)(x_Z(\theta_A, \theta_B) + x_Z(\theta_B, \theta_A)) - b_Z(\theta_A, \theta_B) \geq 0, \quad \forall \theta_A, \theta_B \in \Theta, Z \in \{0, 1\};$$

- *Ex-post per-period budget balance for the designer:*

$$b_Z(\theta_A, \theta_B) \geq p_Z(\theta_A, \theta_B) + p_Z(\theta_B, \theta_A), \quad \forall \theta_A, \theta_B \in \Theta, Z \in \{0, 1\};$$

- *Unit demand constraint:*

$$x_Z(\theta_i, \theta_{-i}) \in [0, 1] \text{ and } x_Z(\theta_i, \theta_{-i}) + x_Z(\theta_{-i}, \theta_i) \leq 1, \quad \forall \theta_i, \theta_{-i} \in \Theta, Z \in \{0, 1\}.$$

Ex-post incentive compatibility requires that it is optimal for dealers to truthfully report their type, for every realization of the opponent's type and conditional on the opponent reporting truthfully.²⁸ Ex-post individual rationality for dealers and customers implies that market participants can walk away from a trade at any time if they find it unprofitable, which is a stronger assumption than the one made in Section 3 (where our solution concept only guaranteed this property for the customers). Ex-post (weak) budget balance states that the designer cannot subsidize trades in any period.²⁹ The unit demand constraint simply ensures that trading probabilities add up to less than one.

²⁸The restriction to stationary mechanisms implies that this condition is independent of continuation payoffs conditional on the state being redrawn in the future.

²⁹This assumption is again made to parallel the setting of Section 3 but it is with loss of generality. We can show that, whenever $\Delta < 1/2$, the designer can strictly increase total welfare if she were allowed to run ex-post budget deficits while still balancing the budget on average. However, this would not change the qualitative features of the optimal mechanism.

An *optimal mechanism* is one that maximizes total discounted probability of trade,

$$\sum_{t=0}^{\infty} \delta^t \mathbb{E}[x_{Z_t}(\theta_{i,t}, \theta_{-i,t}) + x_{Z_t}(\theta_{-i,t}, \theta_{i,t})],$$

over all feasible mechanisms, where the expectation is taken over the sequence $(\theta_{i,t}, \theta_{-i,t}, Z_t)_{t=0}^{\infty}$.

Our first result in this section provides a full characterization of an optimal mechanism.

Proposition 12 (Optimal trading mechanism). *There exists an optimal mechanism in which the implemented outcome—conditional on dealers’ reports—is the same in each period t :*

- *Trade happens with probability one when $(\theta_{i,t}, \theta_{-i,t}) \neq (0, 0)$; otherwise, trade happens with probability $\min\{1, 2\Delta\}$;*
- *Whenever dealers’ types differ, it is the dealer with the lower type who trades;*
- *The dealer who trades receives a payment equal to the other dealer’s type, except when $(\theta_i, \theta_{-i}) = (-1, 0)$, in which case she receives a payment equal to $-1 + \Delta$; the dealer who does not trade receives zero payment;*
- *The customer makes a payment equal to the transfer received by the dealer who trades.*

Expected welfare under an optimal mechanism is higher than welfare in the best-case equilibrium of the pre-trade transparent market (strictly, unless the pre-trade transparent regime achieves the fully efficient outcome).

The first conclusion from Proposition 12 is that there exists an optimal mechanism that takes a *static* form, independent of Z_t . As a result, we will be able to show in the next subsection that the optimal allocation can be implemented using a time-invariant trading protocol—a feature of most real-life trading regimes (and all of the regimes analyzed in Section 3). For intuition, note that the designer faces an infinite sequence of identical two-period problems, whose solution is the same under our restriction to stationary mechanisms.³⁰ Moreover, in each of these two-period problems, the state is fully persistent, in which case it is known that the optimal dynamic mechanism is the repetition of the optimal static one (see, e.g., Baron and Besanko, 1984).

The second conclusion is that the optimal mechanism is fully efficient when $\Delta \geq 1/2$. This is in line with the intuition that, in the presence of adverse selection, efficient outcomes

³⁰This is not true when the mechanism is allowed to be non-stationary. There, the designer can use continuation payoffs to relax dealers’ incentive-compatibility constraints, and therefore the problem that she faces each time the state is redrawn is not static. An example of an implausible non-stationary mechanism that may sometimes improve upon the optimal stationary mechanism is one where, whenever dealers reports coincide, the tie-breaking rule determining who gets to trade depends on what dealers reported in the past.

may become attainable if the gains from trade are sufficiently large. However, it is worth highlighting that the condition $\Delta \geq 1/2$ is not enough to ensure efficient equilibrium trading in any of the trading regimes that we discussed before. When $\Delta < 1/2$, efficient trade is no longer implementable. In any optimal mechanism, trade fails with strictly positive probability in the event that both dealers are uninformed. To sustain trade with probability one when both dealers are uninformed, the transfer given to the low type when her opponent is uninformed would have to be sufficiently high to ensure that the low type does not want to pretend to be uninformed. But when Δ is low, it is not possible to pay the low type such a high transfer without forcing the customer to trade at a loss. Proposition 12 implies that this inefficiency is not a consequence of a suboptimal trading protocol but a fundamental information friction in our trading environment. At the same time, none of the regimes discussed in Section 3 addresses this friction in an optimal way.

The final conclusion from Proposition 12 is that the optimal mechanism resembles a second-price auction (SPA). In an ex-post equilibrium of the SPA, trade happens with the lowest-valuation dealer, and the payment equals the valuation of the losing dealer. The mechanism in Proposition 12 departs from the outcome of the SPA in two ways: Trade happens with interior probability under some contingencies (when both dealers are uninformed and $\Delta < 1/2$), and the payment made by the buyer may lie strictly in between the winning and the losing quote. The first of these departures is a consequence of informational asymmetries between dealers, as described above. The second is a consequence of the fact that the buyer learns from equilibrium quotes. In particular, conditional on the profile of quotes coming from types $(\theta_A, \theta_B) = (-1, 0)$, the customer would be willing to trade at the winning quote but not at the losing quote (since she learns that the value of the asset is low). Hence, the optimal mechanism modifies the price to ensure customer participation.

4.1 An optimal time-invariant trading protocol

Having derived the upper bound on welfare under all possible trading mechanisms, we now turn to constructing a simple trading protocol whose equilibrium implements that optimal outcome. The main result of this subsection says that this is achievable by having dealers compete in every period in a modified version of a descending-clock auction (which is itself a version of the SPA).³¹

Suppose that in every period the trading platform holds the following modified clock auction: The clock starts at some sufficiently high price that is gradually reduced; dealers decide at which price to drop out of the auction. The auction ends when at least one dealer has dropped out, and the trading price equals the current clock price. Ties are broken

³¹Of course, the optimal protocol becomes an ascending-clock auction when the customer is a seller.

uniformly at random. The customer observes dealers' behavior in the auction. When the auction stops, the customer decides between one of three actions: agree, exit, or ask for a new quote. If she agrees, she pays the auction price to the winning dealer and a fee equal to Δ to the platform. If she exits, the game is over. If the customer requests a new quote, the winning dealer makes an additional take-it-or-leave-it offer, which the customer either accepts or rejects. If she accepts, trade takes place at this offer (with no additional fee); if she rejects, there is no trade. There is no post-trade transparency and, in case there is a final take-it-or-leave-it offer after the auction, this is only observed by the customer (it is not observed by the dealer who lost the auction).

Proposition 13 (Indirect implementation of the optimal mechanism). *The allocation and transfers of the optimal mechanism described in Proposition 12 are an ex-post equilibrium outcome of the protocol in which the modified clock auction is run in every period.*

The proof constructs equilibrium strategies of the modified clock auction that implement the optimal outcome. When $Z_t = 0$, dealers quit as soon as the clock announces their estimate of the asset's value. The customer is indifferent between rejecting, accepting or asking for another quote under all on-path outcomes, except when one dealer has a low type and the other one is uninformed. In this event, the customer is not willing to trade at the price prescribed by the auction (0 plus the fee Δ) but she would agree to trade at the low-type's value; thus, the customer requests a new quote from the dealer. Because the dealer has all the bargaining power at this stage of the game, she extracts all the surplus and trade happens with probability one at the price $-1 + \Delta$. These equilibrium strategies do not work when $Z_t = 1$ because in that case dealers must have learned each other's types by observing how the auction unfolded in the previous period (while the period- t customer is uninformed due to post-trade opacity). Nevertheless, we are able to construct the desired mapping between dealers' types and equilibrium outcomes using slightly different strategies.

While the ingredients of the modified clock auction—the SPA pricing rule, trading fees, and post-auction bargaining—are all used on some existing OTC trading platforms (SIFMA, 2016), it is *not* our goal to argue that this particular design should be used in practice. The protocol from Proposition 13 inherits some of the stylized features of our framework (e.g., the fact that the fee is exactly Δ) that would not be present in a richer environment. Instead, our goal is to point out the aspects of the modified clock auction that allow it to outperform the pre-trade transparent trading regime, hoping to identify economic principles that likely hold beyond our simple model.

First, by alleviating the winner's curse for the uninformed type, the SPA pricing rule allows uninformed dealers to quote more aggressively than in the first-price auction. This in

turn makes it less attractive for the low type to pretend to be uninformed, and hence allows to support higher probabilities of trade when both dealers are uninformed. However, the SPA pricing rule makes quotes less attractive for the customer in case one of the dealers is uninformed, which may prevent efficient trade. This issue is resolved by adding the post-auction bargaining stage that allows the winning dealer and the customer to reach mutually agreeable terms of trade. Thus, our results provide a theoretical support for allowing post-auction bargaining on trading platforms in settings with two-sided adverse selection.

The introduction of a fee that must be paid by the customer when trading through the platform serves the purpose of implementing an interior probability of trade conditional on both dealers being uninformed when $\Delta < 1/2$. The fee is not needed when $\Delta \geq 1/2$. Intuitively, the fee is high enough that the customer may sometimes decline trading with the uninformed type, which then reduces the low type’s incentives to pretend to be uninformed. While this intuition is more difficult to apply to real-life trading mechanisms, one conclusion for policy is that the presence of trading fees in adverse selection environments may not necessarily be a sign of inefficiency—eliminating or lowering these fees may lead to worse outcomes due to subtle informational frictions.

Another feature of the trading protocol from Proposition 13 is that the optimal mechanism offers a lot of transparency “within” each trade, but no transparency across trades. This echoes the intuition we emphasized when discussing pre-trade versus full transparency: Once dealers’ information is symmetrized (through seeing each other’s quotes), there is no point in trying to make the customer symmetrically informed as well; in fact, this attempt can backfire as it gives the dealers an incentive to manipulate market beliefs. The clock-auction implementation is one of many ways to symmetrize dealers’ information, and there may be other designs that perform equally well as long as the dealers are able to learn each other’s signals when the auction concludes.

Unlike in Section 3, the insights derived in this section do not require the additional assumption on the dealers’ patience and asset value persistence given in (3.1). In particular, the optimal mechanism can always sustain a fully separating equilibrium, which may be a desirable property from the perspective of the information-aggregation role of prices. However, a caveat of the modified clock auction is that it may have some less efficient equilibria. In Appendix OA.3, we construct an alternative trading regime that achieves unique implementation of the optimal allocation. That trading regime is also a variant of the second price auction with further modifications of the payments, a coarse price grid, and rationing by the platform. These modifications render the design less practical, highlighting a tension between simplicity of trading rules and robustness to equilibrium selection.

5 Concluding remarks

The paper provides a framework for studying several popular trading regimes, as well as the optimal trading protocol, in a model with two-sided adverse selection. The analysis leads to a number of policy implications, such as the ambiguous effects of lowering search costs, superiority of pre-trade transparent markets over post-trade and fully transparent markets, and the potential efficiency-enhancing role of second-price auctions and post-auction bargaining. In order to tractably analyze a variety of trading protocols, we made a number of stylized assumptions about the setting, which could raise concerns about the credibility and robustness of these insights. In this last section, we briefly discuss some extensions as well as limitations of our model.

Multiple dealers. Throughout the paper, we restricted attention to two dealers. OTC markets tend to exhibit a core-periphery trading network (see, for example, [Wang, 2016](#), and the references therein), so assuming that a small number of dealers intermediate most of the trades with customers is realistic. Adding more dealers to the model would complicate the analysis without changing our insights qualitatively.

In the pre-trade opaque market, no-search equilibria would remain the same with N dealers, while search equilibria would have a similar structure: A customer would hope to receive a high quote, and then meet an uninformed dealer. Because the likelihood that some dealer is uninformed conditional on the high state is increasing in N , the search-induced adverse selection problem would become even more severe for uninformed dealers. In the pre-trade transparent market, equilibria would likely become more efficient with large N but adverse selection for uninformed dealers would keep welfare below the efficient level.

Interestingly, even when many dealers are available to provide quotes on OTC trading platforms, auctions tend to involve a small number of dealers. For example, [Riggs et al. \(2020\)](#) document that only about three dealers are contacted in a typical trade on the two largest dealer-to-customer SEFs for credit default swaps. [Wang \(2022\)](#) constructs a theoretical model in which dealers can strategically choose to ignore a request for quote, and shows that, in equilibrium, the customer contacts only two dealers.

Persistence in the asset value. To simplify our analysis, we assumed that the asset value is persistent over at most two periods. We could instead allow it to be persistent over a longer horizon, e.g., by assuming that new information arrives to the market according to a more general stochastic process. However, we believe that such a model would be less tractable without offering qualitatively new insights. In pre-trade transparent regimes, auctions would continue to have the advantageous effect of symmetrizing dealers' information in every period; the remaining benefits of pre-trade transparency—primarily the effect of

direct competition—operate *within* each period. In post-trade transparent regimes, a single informative transaction would lead to efficient trading until new information arrives to the market. Thus, it may seem that allowing for more persistence in the asset value would make post-trade transparency more attractive. That need not be the case, however, because more persistence would also exacerbate dealers’ incentives to manipulate market beliefs (a one-time price manipulation would have a more persistent effect on beliefs of future customers). Finally, our optimal dynamic mechanism was designed to achieve high trade volume within each period and did not rely on intertemporal incentives, so we would expect a similar mechanism to emerge as optimal in environments with higher asset value persistence.

Desirability of post-trade transparency. One of the main insights of our paper is that post-trade transparency has limited value in ensuring market efficiency, assuming the designer has access to other forms of transparency. Post-trade transparency is welfare-dominated by pre-trade transparency, and can only reduce welfare when added to a pre-trade transparent market. The optimal mechanism we constructed in Section 4 also does not feature any post-trade transparency.

The mechanism-design perspective sheds light on the fundamental reasons for the ineffectiveness of post-trade transparency. If the market designer has sufficient control over the market protocol, then post-trade transparency can never be helpful as a consequence of the *inscrutability principle* (Myerson, 1983). The inscrutability principle states, roughly, that if a designer wants to implement a certain outcome, then she cannot gain by disclosing her information (in this case, about past trades) to the agents. The reason is that disclosing information only increases the number of incentive constraints that the designer must satisfy; if all constraints hold conditional on any disclosure, then they must also hold in expectation, that is, when information is withheld. In the language of indirect trading mechanisms, disclosure of endogenous variables—such as the price—creates incentives for agents to manipulate them for their private benefit. Consequently, as long as the designer can choose the trading protocol within each period, disclosing information about past trades only makes the design problem more difficult.

That being said, post-trade transparency may be desirable when the designer has limited control over the market protocol. For example, implementing pre-trade transparency may be politically costly in some markets due to the concentrated interest of the dealers. Then, by Proposition 6, introducing post-trade transparency can increase welfare relative to a fully opaque market. This result is in line with support for post-trade transparency offered by, for example, Duffie et al. (2017) and Back et al. (2020). Relative to these papers, Proposition 6 shows that post-trade transparency may be welfare-enhancing despite the negative effect it

has through belief manipulation.³²

Post-trade transparency may have benefits not captured by our model. First, [Duffie et al. \(2017\)](#) show that post-trade transparency (introduction of a benchmark) allows traders to make more efficient entry decisions (primarily by avoiding costly entry when gains from trade are small).³³ Second, if investors acquire information before trading, post-trade transparency could help reduce the costs of information acquisition. Third, post-trade transparency may be an important source of information about the market for actors not directly interested in trading, such as regulators or real-economy decisions makers. While this last benefit could be substantial, we note that our critique of post-trade transparency applies only to *immediate* public disclosure of trade data (the belief manipulation effect is present only when dealers’ information is relevant for transactions following the disclosure). Our results are fully consistent with the idea that it can be beneficial to disclose trade data with a (small) time lag—and indeed, our focus on separating equilibria means that such “lagged” post-trade transparency would have the same social value in all the regimes we have studied.

Information acquisition. Throughout the paper, we treated the information structure as exogenous. It would be conceptually straightforward to endogenize the information structure by letting dealers choose the probability q of being informed optimally at an ex-ante stage subject to a cost. It is well known that more competitive market protocols—such as our pre-trade transparent market—may have adverse effects on the incentives of market participants to acquire information.³⁴ Thus, we may expect that dealers would acquire more information under market protocols that give them some degree of market power. Since inefficiencies in our model vanish at both extremes—with fully informed as well as fully uninformed dealers—the effects of information acquisition on welfare are not ex-ante obvious. We leave this direction for future research.

Inventory management and the inter-dealer market. Our model focuses on the dealer-to-customer segment of OTC markets. However, a large fraction of trades happen between dealers trying to rebalance their positions—an aspect that our framework does not capture due to the simplifying assumption of unit trade and constant marginal values. Before the introduction of TRACE, dealers expressed a concern that too much transparency would hurt their ability to intermediate trades by raising their price impact in the inter-dealer

³²In [Duffie et al. \(2017\)](#), the belief manipulation effect is absent because information disclosed to the market is exogenous; in the model of [Back et al. \(2020\)](#), the potential negative effect of belief manipulation is neutralized by the assumption that trade always takes place in the downstream market.

³³Because pre-trade transparency shifts more rents towards investors, it is not clear which form of transparency would lead to more efficient entry in a model with costly entry and stochastic gains from trade.

³⁴[Grossman and Stiglitz \(1980\)](#) is perhaps the most famous example, and [Baldauf and Mollner \(2020\)](#) is a recent example in the context of a market microstructure model.

market (Asquith et al., 2019). And indeed, dealers might prefer to use non-transparent mechanisms such as dark pools and size-discovery sessions to avoid price impact (see, for example, Zhu, 2012 and Duffie and Zhu, 2017). Of course, dealers’ opposition to transparent mechanisms does not imply that transparency is not welfare-enhancing overall (for example, dealers in our model would lose from a transition from the opaque to the pre-trade transparent market). The question is whether dealers’ reluctance to take on large positions in the face of a transparent inter-dealer market could reduce market efficiency.

Answering this question would require a model with non-linear dealer values and an inter-dealer market, such as the one proposed by Babus and Parlato (2022)—an extension of the linear-quadratic double-auction model of Rostek and Weretka (2012). Unfortunately, the tractability of the double-auction model (and—to the best of our knowledge—most financial market models featuring a meaningful notion of trade size) is limited to linear equilibria and Gaussian information structures that are typically not preserved under interesting modifications of market transparency.³⁵ Thus, understanding optimal market transparency in the presence of inter-dealer markets remains an important open problem.

References

- ANTILL, S. AND D. DUFFIE (2020): “Augmenting Markets with Mechanisms,” The Review of Economic Studies, 88, 1665–1719.
- AQUILINA, M., G. IBIKUNLE, V. MOLLICA, AND T. STEFFEN (2022): “The visible hand: benchmarks, regulation, and liquidity,” Journal of Financial Markets, 100734.
- ASQUITH, P., T. COVERT, AND P. A. PATHAK (2019): “The Effects of Mandatory Transparency in Financial Market Design: Evidence from the Corporate Bond Market,” Available at SSRN 2320623.
- ASRIYAN, V., W. FUCHS, AND B. GREEN (2017): “Information Spillovers in Asset Markets with Correlated Values,” American Economic Review, 107, 2007–40.
- BABUS, A. AND C. PARLATORE (2022): “Strategic fragmented markets,” Journal of Financial Economics, 145, 876–908.
- BACK, K., R. LIU, AND A. TEGUIA (2020): “Signaling in OTC Markets: Benefits and Costs of Transparency,” Journal of Financial and Quantitative Analysis, 55, 47–75.
- BAGWELL, K. AND G. RAMEY (1991): “Oligopoly limit pricing,” The Rand Journal of Economics, 155–172.
- BALDAUF, M. AND J. MOLLNER (2020): “High-Frequency Trading and Market Performance,” The Journal of Finance, 75, 1495–1526.

³⁵The only exception we are aware of is Ollar et al. (2021) but even that paper can only analyze a special class of transparency regimes.

- BANKS, J. S. AND J. SOBEL (1987): “Equilibrium selection in signaling games,” Econometrica: Journal of the Econometric Society, 647–661.
- BARON, D. P. AND D. BESANKO (1984): “Regulation and information in a continuing relationship,” Information Economics and policy, 1, 267–302.
- BARSANETTI, B. AND B. CAMARGO (2022): “Signaling in dynamic markets with adverse selection,” Journal of Economic Theory, 206, 105558.
- BERGEMANN, D. AND J. HÖRNER (2018): “Should first-price auctions be transparent?” American Economic Journal: Microeconomics, 10, 177–218.
- BESSEMBINDER, H. AND W. MAXWELL (2008): “Markets: Transparency and the corporate bond market,” Journal of Economic Perspectives, 22, 217–234.
- BLOUIN, M. R. AND R. SERRANO (2001): “A decentralized market with common values uncertainty: Non-steady states,” The Review of Economic Studies, 68, 323–346.
- BUDISH, E., P. CRAMTON, AND J. SHIM (2015): “The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response *,” The Quarterly Journal of Economics, 130, 1547–1621.
- CHAVES, I. N. (2020): “Privacy in bargaining: The case of endogenous entry,” Available at SSRN 3420766.
- CHO, I.-K. AND D. M. KREPS (1987): “Signaling games and stable equilibria,” The Quarterly Journal of Economics, 102, 179–221.
- CRÉMER, J. AND R. P. MCLEAN (1985): “Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent,” Econometrica, 53, 345–362.
- DALEY, B. AND B. GREEN (2012): “Waiting for News in the Market for Lemons,” Econometrica, 80, 1433–1504.
- DIAMOND, P. A. (1971): “A model of price adjustment,” Journal of Economic Theory, 3, 156–168.
- DUFFIE, D. (2012): Dark Markets: Asset Pricing and Information Transmission in Over-the-Counter Markets, Princeton University Press.
- DUFFIE, D., P. DWORCZAK, AND H. ZHU (2017): “Benchmarks in Search Markets,” The Journal of Finance, 72, 1983–2044.
- DUFFIE, D., N. GÂRLEANU, AND L. H. PEDERSEN (2005): “Over-the-Counter Markets,” Econometrica, 73, 1815–1847.
- DUFFIE, D. AND J. C. STEIN (2015): “Reforming LIBOR and Other Financial Market Benchmarks,” Journal of Economic Perspectives, 29, 191–212.
- DUFFIE, D. AND H. ZHU (2017): “Size Discovery,” The Review of Financial Studies, 30, 1095–1150.
- DWORCZAK, P. (2020): “Mechanism Design With Aftermarkets: Cutoff Mechanisms,” Econometrica, 88, 2629–2661.

- DYE, R. A. (1985): “Disclosure of Nonproprietary Information,” Journal of Accounting Research, 23, 123–145.
- EDWARDS, A. K., L. E. HARRIS, AND M. S. PIWOWAR (2007): “Corporate bond market transaction costs and transparency,” The Journal of Finance, 62, 1421–1451.
- FAMA, E. F. AND M. H. MILLER (1972): The Theory of Finance, Holt, Rinehart, and Winston, New York.
- FUCHS, W., P. GOTTARDI, AND H. MOREIRA (2022): “Time Trumps Quantity in the Market for Lemons,” Working Paper.
- FUCHS, W., A. ÖRY, AND A. SKRZYPACZ (2016): “Transparency and distressed sales under asymmetric information,” Theoretical Economics, 11, 1103–1144.
- GLEBKIN, S., B. Z. YUESHEN, AND J. SHEN (2023): “Simultaneous multilateral search,” The Review of Financial Studies, 36, 571–614.
- GLOSTEN, L. R. AND P. R. MILGROM (1985): “Bid, ask and transaction prices in a specialist market with heterogeneously informed traders,” Journal of Financial Economics, 14, 71–100.
- GROSSMAN, S. J. AND J. E. STIGLITZ (1980): “On the Impossibility of Informationally Efficient Markets,” The American Economic Review, 70, 393–408.
- GUERRIERI, V. AND R. SHIMER (2014): “Dynamic adverse selection: A theory of illiquidity, fire sales, and flight to quality,” American Economic Review, 104, 1875–1908.
- GUERRIERI, V., R. SHIMER, AND R. WRIGHT (2010): “Adverse selection in competitive search equilibrium,” Econometrica, 78, 1823–1862.
- GUL, F. AND W. PESENDORFER (2012): “The war of information,” The Review of Economic Studies, 79, 707–734.
- HALAC, M. AND I. KREMER (2020): “Experimenting with career concerns,” American Economic Journal: Microeconomics, 12, 260–88.
- HENDERSHOTT, T. AND A. MADHAVAN (2015): “Click or call? Auction versus search in the over-the-counter market,” The Journal of Finance, 70, 419–447.
- HÖRNER, J. AND N. VIEILLE (2009): “Public vs. private offers in the market for lemons,” Econometrica, 77, 29–69.
- HOU, D. AND D. SKEIE (2013): “LIBOR: Origins, Economics, Crisis, Scandal and Reform,” in The New Palgrave Dictionary of Economics, Online Edition, ed. by S. N. Durlauf and L. E. Blume, Palgrave Macmillan.
- INDERST, R. (2005): “Matching markets with adverse selection,” Journal of Economic Theory, 121, 145–166.
- KAKHBOD, A. AND F. SONG (2020): “Dynamic price discovery: Transparency vs. information design,” Games and Economic Behavior, 122, 203–232.
- (2022): “Public vs. Private offers with informed and forward-looking dealers,” Working Paper.

- KAYA, A. AND K. KIM (2018): “Trading dynamics with private buyer signals in the market for lemons,” The Review of Economic Studies, 85, 2318–2352.
- KAYA, A. AND S. ROY (2022): “Market screening with limited records,” Games and economic behavior, 132, 106–132.
- KIM, K. (2017): “Information about sellers’ past behavior in the market for lemons,” Journal of Economic Theory, 169, 365–399.
- KYLE, A. S. (1985): “Continuous Auctions and Insider Trading,” Econometrica, 53, 1315–1335.
- LAUERMANN, S. AND A. WOLINSKY (2016): “Search with adverse selection,” Econometrica, 84, 243–315.
- LESTER, B., A. SHOURIDEH, V. VENKATESWARAN, AND A. ZETLIN-JONES (2019): “Screening and adverse selection in frictional markets,” Journal of Political Economy, 127, 338–377.
- MILGROM, P. AND J. MOLLNER (2021): “Extended proper equilibrium,” Journal of Economic Theory, 194, 105258.
- MYERSON, R. B. (1983): “Mechanism Design by an Informed Principal,” Econometrica, 51, 1767–1797.
- OLLAR, M., M. J. ROSTEK, AND J. H. YOON (2021): “Privacy in Markets,” Available at SSRN: <https://ssrn.com/abstract=3071374>.
- OSBORNE, M. J. AND A. RUBINSTEIN (1990): Bargaining and markets, Academic press.
- RIGGS, L., E. ONUR, D. REIFFEN, AND H. ZHU (2020): “Swap trading after Dodd-Frank: Evidence from index CDS,” Journal of Financial Economics, 137, 857–886.
- ROSTEK, M. AND M. WERETKA (2012): “Price Inference in Small Markets,” Econometrica, 80, 687–711.
- ROTHSCHILD, M. AND J. STIGLITZ (1976): “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” The Quarterly Journal of Economics, 90, 629–649.
- SIFMA (2016): “[SIFMA Electronic Bond Trading Report: US Corporate and Municipal Securities](#),” Report.
- VIDA, P. AND T. HONRYO (2021): “Strategic stability of equilibria in multi-sender signaling games,” Games and Economic Behavior, 127, 102–112.
- WANG, C. (2016): “Core-Periphery Trading Networks,” Available at SSRN: <https://ssrn.com/abstract=2747117>.
- (2022): “The Limits of Multi-Dealer Platforms,” Available at SSRN: <https://ssrn.com/abstract=4199500>.
- WOLINSKY, A. (1990): “Information revelation in a market with pairwise meetings,” Econometrica: Journal of the Econometric Society, 1–23.
- ZHU, H. (2012): “Finding a Good Price in Opaque Over-the-Counter Markets,” The Review of Financial Studies, 25, 1255–1285.

A Additional comments on related literature

In this appendix, we explain in more detail the relationship between our paper and the literature on search and bargaining under adverse selection, focusing especially on papers that studied the effects of information disclosure.

A classic question studied within the literature on search under adverse selection concerns the issue of information aggregation of prices when trading is decentralized (Wolinsky, 1990, Blouin and Serrano, 2001). In these papers, there are two sides of the market, each composed by a continuum of agents (e.g., buyers and sellers), who have private information about the common value of the good. Trade happens through repeated pairwise meetings, and the outcome of previous encounters is unobservable to other market participants. Their key finding is that, as the market becomes approximately frictionless, prices fail to fully reveal the true value of the good. Blouin and Serrano (2001) also take into account allocative efficiency, and find that the equilibrium outcome may only be inefficient when both sides of the market have private information. Unlike them, we also find inefficiency under one-sided private information when the search friction is negligible. Even though there are several other differences in the modeling assumptions used in those papers relative to this one, a key factor driving the discrepancy in the results comes from the trading procedure. In their setting, the two sides simultaneously announce their trading positions, which implies that inefficiency stems from adverse selection, and in particular from uninformed traders taking on a tough position because they fear trading at a loss. In our model, the informed party makes the first move, giving rise to an additional source of inefficiency due to signaling incentives.³⁶ An alternative approach, pursued by Inderst (2005), Guerrieri et al. (2010), and Lester et al. (2019), augment the competitive screening environment from Rothschild and Stiglitz (1976) by introducing directed search and characterize the set of equilibrium contracts. Those papers also differ from ours in that it is the uninformed party (the competing principals) who propose the terms of trade, and the informed party (the agents) who conduct the search.

Lauermann and Wolinsky (2016) study a model of one-sided sequential search under adverse selection in which, unlike in our setting, the informed party is the one who searches. Despite differences between the models, they identify a similar force as the one underlying some of our welfare results: Search trading regimes are in disadvantage relative to auction-type ones due to a strong form of winner’s curse, whereby the event of being contacted for a quote leads sellers to infer that the value of the asset is high and thus price more aggressively. In their setting, this inference by the seller arises due to exogenous private

³⁶Barsanetti and Camargo (2022) study a version of the model in Wolinsky (1990) where the informed seller proposes the price, and like us find that signaling through prices gives rise to inefficiency, even when the buyer does not have private information.

information possessed by the buyer, whereas in ours it is the result of the searching party gradually becoming more informed by observing the informed seller's price offers.

We also connect to the literature that studies the effects of transparency and information disclosure in dynamic adverse selection settings. Like us, a series of these papers study a long-run privately informed seller that trades with a sequence of uninformed short-lived buyers. In Hörner and Vieille (2009) and Fuchs et al. (2016), the seller possesses a single unit of an indivisible good, and inefficiency may arise in equilibrium due to the high quality seller delaying its sale in order to signal her type. Both papers find that public disclosure of previously rejected offers leads to lower equilibrium welfare. The reason is that, under public offers, the high quality seller may use rejecting a high price as an additional tool to signal her type. Kim (2017), who studies the effect of observability of the seller's time on the market, arrives to a similar conclusion.³⁷ Kaya and Roy (2022) study the role of the availability of past trading records in a setting with repeated sales, and find that welfare is non-monotone in record length. In Kaya and Kim (2018), the buyer observes a private signal about the seller's type. As a result of two-sided private information, transaction delay may now become a signal of the quality being low, due to pessimistic privately informed buyers refusing to trade. In the multi-unit environment studied by Fuchs et al. (2022), the seller may resort to both delay and a lower volume of trade in order to signal high quality, although they find that only the former is used in equilibrium. In their model, public disclosure of past transactions does not impact welfare in the continuous time limit, although it does have (ambiguous) welfare effects when trade happens in discrete time.

Our model differs from these papers in several dimensions. First, in our setting, trade happens repeatedly over time and it is the informed party (the dealer) who proposes the terms of trade. Signaling of high quality by the dealers will therefore take place through prices. Because of this, inefficiency arises in our model through interior probabilities of trade in order to deter the lower types from mimicking a high type who trades at a higher price, instead of through delay of a one-time transaction. Moreover, the aforementioned papers feature a single seller who, in every period, receives an offer from a new buyer, which the seller chooses whether or not to accept. Thus, the seller effectively conducts sequential search (with the discount factor representing the search friction). In our setting, there is more than one seller, and we study different trading regimes that differ both in the extent of public disclosures of information, and in the amount of competition between sellers.

³⁷Chaves (2020) shows that the effects of post-trade transparency can be very different in a setting in which there is bargaining between two long-lived agents and there is endogenous entry to the market.

B Proofs

Preliminaries: Two-period interpretation of the model

In order to simplify the exposition, we will make use of the convenient fact that, under our assumptions, the set of stationary equilibrium strategies in the trading regimes that we study coincide with the equilibrium strategies of the following two-period game. In period zero, v is drawn, dealers privately observe their signals, and the trading game is played. In period one, dealers play again the trading game, with the value of v remaining fixed and dealers observing no further signal realizations. Dealers discount the future by a factor of $\delta\gamma$.

The equivalence holds because, under the stationarity restriction, dealers' strategies only depend on their private information and the realization of $Z_t \in \{0, 1\}$. Since all variables are i.i.d. redrawn when $Z_t = 0$, the game “resets” in this event, in the sense that the resulting continuation game is history-independent. Moreover, in that game, dealers only contemplate their current payoffs and their continuation payoffs conditional on the state remaining fully persistent in the following period, and they will effectively “discount” this continuation payoff by a factor of $\delta\gamma$ (the discount factor times the probability of reaching this stage of the game). When $Z_t = 1$, players know with certainty that the game will reset in the following period, and therefore behave as if this was the final period of the game. As a result, stationary equilibrium strategies in the infinite-horizon game when $Z_t = 0$ and $Z_t = 1$ are the same as equilibrium strategies in the two-period game in period zero and period one, respectively.

In addition, the following lemma shows that, regardless of the trading regime, discounted expected equilibrium welfare in the infinite-horizon game is proportional to its counterpart in the two-period game. Thus, the welfare rankings across trading regimes are the same in the two games.

Lemma 1. *Equilibrium welfare under protocol r is*

$$\Lambda_r = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \mathbb{E}[\Lambda_r(Z_t)] = \frac{\Lambda_r(0) + \delta\gamma\Lambda_r(1)}{1 + \delta\gamma},$$

where $\Lambda_r(Z)$, $Z \in \{0, 1\}$, are equilibrium probabilities of trade conditional on $Z_t = Z$.

Proof. Let $p_t = (\Pr(Z_t = 0), \Pr(Z_t = 1))$, be the ex ante distribution of Z_t . Since Z_t follows a Markov chain, we can write

$$p_t = p_0 \times P^t,$$

where $p_0 = (1, 0)$ is the distribution of Z_0 and

$$P \equiv \begin{pmatrix} 1 - \gamma & \gamma \\ 1 & 0 \end{pmatrix}$$

is the Markov transition matrix of Z_t .

P can be decomposed as

$$P = A \begin{pmatrix} 1 & 0 \\ 0 & -\gamma \end{pmatrix} A^{-1}, \quad A = \begin{pmatrix} 1 & -\gamma \\ 1 & 1 \end{pmatrix}.$$

Hence,

$$p_t = p_0 \times A \begin{pmatrix} 1 & 0 \\ 0 & (-\gamma)^t \end{pmatrix} A^{-1} = \left(\frac{1 + \gamma(-\gamma)^t}{1 + \gamma}, \frac{\gamma(1 - (-\gamma)^t)}{1 + \gamma} \right).$$

Given $\Lambda_r(Z)$, the discounted expectation can be written as

$$\sum_{t=0}^{\infty} \delta^t p_t \times \begin{pmatrix} \Lambda_r(0) \\ \Lambda_r(1) \end{pmatrix} = \frac{\Lambda_r(0) + \delta\gamma\Lambda_r(1)}{(1 - \delta)(1 + \delta\gamma)}.$$

□

Given the weights that arise from Lemma 1, we define

$$\alpha \equiv \frac{1}{1 + \delta\gamma}. \tag{B.1}$$

With this equivalence in mind, throughout the Appendix, we refer to realizations of $Z_t = z \in \{0, 1\}$ simply as period z .

Proof of Proposition 1

We start by establishing an auxiliary result that will be useful when describing equilibria in the regimes with sequential search. To do so, we consider an auxiliary *one-shot search game*, which is the game that arises in a one-period version of the opaque market. Let b_σ^n be any price in the support of the equilibrium quoting strategy of a dealer with initial type $\sigma \in \{l, u, h\}$ at the n th time she is contacted by the customer, and let $\underline{b}_\sigma^n, \bar{b}_\sigma^n$ be the respective infimum and supremum of the support. Implicitly, this quote is allowed to depend on the $n - 1$ prices that the dealer quoted in the past.

Lemma 2. *In all equilibria of the one-shot search game, on the equilibrium path: (i) no dealer gets contacted more than once, (ii) conditional on being contacted, the low type quotes $-1 + \Delta$ and trades with probability one, and (iii) if there is no search in equilibrium, then $b_1^u \geq \Delta$ and $b_1^h \geq 1 + \Delta$. Moreover, there is $\bar{s} > 0$ such that, if $s < \bar{s}$: (iv) the customer may only search in equilibrium after receiving a quote equal to b_h^1 , and (v) if there is search in equilibrium, $b_u^1 = \Delta$ and $b_h^1 = 1 + \Delta$, and the customer agrees to trade with probability one after observing a pair of quotes (b_h^1, b_h^1) .*

Proof. Part (i). Suppose by contradiction that, at the n th time that she gets contacted, dealer i with type σ expects to be contacted again by the customer with strictly positive probability. Let v'_b denote the buyer's belief about v at the $(n + 1)$ th time that she contacts dealer i (which depends on the buyer's private history). The fact that the buyer is willing to contact dealer i for the $(n + 1)$ th time implies that $v'_b + \Delta - \underline{b}_\sigma^{n+1} - s > 0$; otherwise, the customer would strictly prefer to exit the market over contacting dealer i again. It also implies that $v'_b + \Delta - \underline{b}_\sigma^{n+1}$ is weakly greater than the expected payoff from contacting dealer $j \neq i$ instead. But then, since $\underline{b}_\sigma^{n+1} < v'_b + \Delta$, dealer i can strictly benefit by deviating at the $(n + 1)$ th contact to $\tilde{b}_\sigma^{n+1} = \underline{b}_\sigma^{n+1} + \epsilon$, with $0 < \epsilon < s$. By monotonicity of beliefs, this higher quote may only weakly improve the buyer's estimate of the asset's value. Thus, accepting this offer strictly dominates exiting the market or paying the search cost to contact dealer j . This implies that dealer i can trade at a higher price with probability one, a contradiction.

Part (ii). By part (i), it suffices to focus on b_l^1 . First, note that $\bar{b}_l^1 \leq -1 + \Delta$; otherwise, a quote equal to \bar{b}_l^1 is accepted by the customer with zero probability, and the low-type dealer could do strictly better by offering $-1 + \Delta - \epsilon$ (which is always accepted by the customer) provided that $\epsilon > 0$ is small enough. Second, we show that $\underline{b}_l^1 \geq -1 + \Delta$. Towards a contradiction, suppose that $\underline{b}_l^1 < -1 + \Delta$. We claim that the low-type dealer i can trade with probability one at a quote higher than \underline{b}_l^1 , thus yielding a profitable deviation. Specifically, consider dealer i 's payoff from deviating to $\tilde{b} = \underline{b}_l^1 + \epsilon$ where $\epsilon > 0$ is arbitrarily small. Consider first the case in which i is the customer's first contact. Then, the customer has a strict incentive to accept when quoted \underline{b}_l^1 . This is because, if she searches after this quote, she expects that the other dealer will quote a price weakly greater than \underline{b}_l^1 with probability one, and a price strictly greater than \underline{b}_l^1 with probability $1 - q > 0$ (since the uninformed type will never quote below 0). Thus, searching after \underline{b}_l^1 is strictly suboptimal; and so is exiting the market given that $\underline{b}_l^1 < -1 + \Delta$. This implies that the first-contact customer will also have a strict incentive to accept \tilde{b} if we take ϵ to be small enough. On the other hand, if dealer i is the customer's second contact, we know by part (i) that she must accept \underline{b}_l^1 with probability one. Moreover, it cannot be that she is indifferent between accepting and searching at this node; for if she were, the dealer whom she is contacting for a second time

(off-path) must be quoting strictly less than \underline{b}_i^1 and can do better by increasing her quote by the same argument as in part (i). As a result, the second-contact customer also has a strict incentive to accept \underline{b}_i^1 , and thus the same is true of \tilde{b} . So, regardless of it being the first or the second contact by the customer, dealer i can secure that the customer accepts $\tilde{b} > \underline{b}_i^1$ with probability one, a contradiction.

Summarizing, we have $\underline{b}_i^1 = \bar{b}_i^1 = -1 + \Delta$. Finally, the fact that $b_u^1, b_h^1 \geq -1 + \Delta$ with probability one,³⁸ implies that a quote of $-1 + \Delta$ must be accepted by the customer with probability one, since otherwise the low-type dealer could offer $-1 + \Delta - \epsilon$ which the customer would strictly want to accept.

Part (iii). If there is no search in equilibrium and the uninformed type quotes $0 \leq b_u^1 < \Delta$, the customer agrees to trade for sure, since she strictly prefers trading at that price to exiting the market. But this gives rise to a profitable deviation for the low type—by quoting b_u^1 , she can trade for sure (conditional on being contacted) at a strictly higher price. The same argument rules out $b_h^1 < 1 + \Delta$.

Part (iv). First, given what we established in part (ii), the customer clearly does not want to search following a quote from a low-type dealer. Suppose that the customer searches after observing b_u^1 ; we will derive a contradiction when the search cost s is small enough. Since $b_i^1 = -1 + \Delta$, a necessary condition for searching to be optimal for the customer is that $\Delta - b_u^1 \leq (1 - q) \max\{\Delta - \underline{b}_u^1, 0\} + \frac{q}{2} \max\{1 + \Delta - \underline{b}_h^1, 0\} - s$. In particular, it has to be the case that either $\underline{b}_u^1 < \Delta$ or that $\underline{b}_h^1 < 1 + \Delta$. Suppose first that $\underline{b}_h^1 \geq 1 + \Delta$ and $\underline{b}_u^1 < \Delta$. Then, \underline{b}_u^1 is the lowest quote at which the customer can trade at a profit, and therefore she must accept this quote with probability one independent of whether this is the customer's first or second contact. But then, the uninformed type can do strictly better by slightly increasing her quote to $\underline{b}_u^1 + \epsilon$ which the customer would still have a strict incentive to accept.

Next, suppose that $\underline{b}_h^1 < 1 + \Delta$. This implies that the following must be true:

1. \underline{b}_h^1 is such that the customer is indifferent between accepting and searching after observing a first quote equal to \underline{b}_h^1 . Otherwise, the high type could profitably deviate by slightly increasing her quote, which would not affect the “first-contact” customer's decision, and would still get the “second-contact” customer to agree to trade with probability one (by belief monotonicity and $\underline{b}_h^1 < 1 + \Delta$, she would strictly want to accept).
2. The support of the high-type's quoting strategy is a subset of $\{\underline{b}_h^1, 1 + \Delta\}$. This is because, by the previous step, any quote above \underline{b}_h^1 may only trade against a second-contact customer, and is therefore strictly dominated by quoting $1 + \Delta$.

³⁸The same argument as the one in the previous paragraph rules out $\underline{b}_\sigma^1 < -1 + \Delta$ for $\sigma = u, h$.

3. $\bar{b}_u^1 \leq \Delta$. Otherwise, the uninformed type who quotes \bar{b}_u^1 may only trade when facing a second-contact customer. If $\bar{b}_u^1 < \underline{b}_h^1$, the uninformed type can do strictly better by deviating to \underline{b}_h^1 which still leads to trade with probability one with a second-contact customer. If $\bar{b}_u^1 > \underline{b}_h^1 \geq 1$, the uninformed type may only trade if the first contact by the customer was a high type. Hence, it must be that $\bar{b}_u^1 = 1 + \Delta$ and the customer accepts with probability one; otherwise, the uninformed type could slightly increase her quote and still get the customer to trade. But the fact that the high-type weakly prefers to quote \underline{b}_h^1 over $1 + \Delta$, implies that the uninformed type (whose estimate of the asset's value is strictly less than 1) *strictly* prefers the former over the latter, a contradiction.
4. b_u^1 is such that the customer is indifferent between searching and accepting after receiving a first quote equal to b_u^1 . If $b_u^1 < \Delta$, this follows from the same argument as in point 1. If $b_u^1 = \Delta$ and the customer strictly prefers searching over accepting, the quote b_u^1 only leads to trade against a second-contact customer. But then the uninformed type can trade with the same probability at a strictly higher price by deviating to \underline{b}_h^1 , a contradiction.
5. $b_u^1 = \underline{b}_u^1 = \bar{b}_u^1$. Suppose first that $b_u^1 > \underline{b}_u^1$. By monotonicity of beliefs and point 4, the uninformed dealer can then profit from slightly lowering b_u^1 and getting the first-contact customer to accept with probability one. Therefore, $b_u^1 = \underline{b}_u^1$, that is, there can only be search following the lowest quote in the support. Next, suppose that $\bar{b}_u^1 > \underline{b}_u^1$. Again, by step 4, the first-contact customer will strictly want to reject \bar{b}_u^1 , so this quote leads to trade only against the second-contact customer. But then the uninformed dealer can do strictly better by quoting \underline{b}_h^1 instead, a contradiction.

Combining the above results, the following condition ensuring that the customer is indifferent between searching and accepting after both \underline{b}_h^1 and \underline{b}_u^1 must be satisfied:

$$1 + \Delta - \underline{b}_h^1 = (1 - q)(1 + \Delta - \underline{b}_u^1) + qp_h(1 + \Delta - \underline{b}_h^1) - s, \quad (\text{B.2})$$

$$\Delta - \underline{b}_u^1 = (1 - q)(\Delta - \underline{b}_u^1) + \frac{q}{2}p_h(1 + \Delta - \underline{b}_h^1) - s, \quad (\text{B.3})$$

where $p_h \in (0, 1]$ is the probability that the high-type quotes \underline{b}_h^1 . Note that, when s is small enough, the first condition can be satisfied only if $p_h < 1$. This implies that the following condition ensuring that the high type is indifferent between \underline{b}_h^1 and $1 + \Delta$ must be satisfied:

$$\begin{aligned} (\rho_h + (1 - q)(1 - \rho_u) + q(p_h(1 - \rho_h) + 1 - p_h))(\underline{b}_h^1 - 1) \\ = ((1 - q)(1 - \rho_u) + q(p_h(1 - \rho_h) + 1 - p_h))\Delta, \end{aligned} \quad (\text{B.4})$$

where ρ_h and ρ_u are the respective probabilities that the first-contact customer agrees to

trade when she receives the quotes \underline{b}_h^1 and \underline{b}_u^1 .

For (B.2),(B.3) to admit a solution where $\underline{b}_h^1 \geq 1$, we need that $\Delta > 1 - q$, so we can assume that this holds for the remainder of this proof.

Solving (B.2), (B.3), (B.4), and taking the limit as s goes to zero gives

$$\underline{b}_u^1 = \Delta - \frac{\Delta\rho_h - (1-q)(1+\rho_h - \rho_u(1-q))}{1 - \rho_u + q(1 - \rho_h) + \rho_h + q^2\rho_u} < \Delta \iff \rho_h > \frac{(1-q)(1 - (1-q)\rho_u)}{\Delta - (1-q)} \equiv \underline{\rho}_h.$$

Now, consider the constraint preventing the uninformed type from deviating to \underline{b}_h^1 :

$$\begin{aligned} & \rho_u \underline{b}_u^1 + (1-q)(1 - \rho_u) \underline{b}_u^1 + \frac{q}{2}(p_h(1 - \rho_h) + 1 - p_h)(\underline{b}_u^1 - 1) \geq \\ & \rho_h \underline{b}_h^1 + (1-q)(1 - \rho_u) \underline{b}_h^1 + \frac{q}{2}(p_h(1 - \rho_h) + 1 - p_h)(\underline{b}_h^1 - 1) \iff \\ & \rho_u \underline{b}_u^1 - \rho_h \underline{b}_h^1 + \left[(1-q)(1 - \rho_u) + \frac{q}{2}(p_h(1 - \rho_h) + 1 - p_h) \right] (\underline{b}_u^1 - \underline{b}_h^1) \geq 0. \end{aligned}$$

On the other hand, plugging in $\rho_h > \underline{\rho}_h$, we have:

$$\begin{aligned} & \rho_u \underline{b}_u^1 - \rho_h \underline{b}_h^1 + \left[(1-q)(1 - \rho_u) + \frac{q}{2}(p_h(1 - \rho_h) + 1 - p_h) \right] (\underline{b}_u^1 - \underline{b}_h^1) < \\ & [2(\Delta - 1 + q)]^{-1} [(\underline{b}_u^1 - \underline{b}_h^1)(2\Delta(1 - q\rho_u) + q(1 - q) + q^2p_h(1 - q\rho_u) + \Delta q + p_h q(1 - \rho_u)) \\ & \quad - 2\underline{b}_u^1(1 - q + (1 - p_h)q^2) - 2\rho_u(\Delta \underline{b}_h^1 - q\underline{b}_u^1 - q^2p_h \underline{b}_h^1)] < 0, \end{aligned}$$

where the last inequality comes from $\underline{b}_u^1 \leq \Delta < 1 \leq \underline{b}_h^1$. This yields a contradiction.

Part (v). Suppose that there is search in equilibrium. If $b_u^1 < \Delta$, given that we have ruled out search following b_u^1 , the customer must agree to trade with probability one after observing b_u^1 . But if this were the case, the low-type dealer would want to deviate to b_u^1 , which allows her to trade at a higher price with probability one. So $b_u^1 \geq \Delta$.

Moreover, if $b_u^1 > \Delta$ and there is search in equilibrium, the uninformed dealer who quotes b_u^1 can only trade if the customer contacted a high-type dealer first and learned that $v = 1$. If this is the case, the uninformed dealer can get the customer to agree to trade at any quote below $1 + \Delta$. So it must be that $b_u^1 = 1 + \Delta$ and this quote is accepted by the second-contact customer with probability one. But then the customer would not find it profitable to search in the first place, a contradiction. Thus, we have $b_u^1 = \Delta$ whenever the equilibrium has search. A similar argument to the one in part (iv) shows that, if there is search in equilibrium, it must be that $b_h^1 = 1 + \Delta$. Observe that, when s is small enough, the first-contact customer will strictly prefer searching over accepting $1 + \Delta$ (since she trades at a profit if she meets an uninformed dealer), so she will reject this first quote with probability one.

After observing a pair of quotes $(1 + \Delta, 1 + \Delta)$, the customer is indifferent between trading

and exiting the market. However, it must be that she agrees to trade for sure; otherwise, the high-type dealer could deviate to $1 + \Delta - \epsilon$ for a small $\epsilon > 0$. This deviation would be profitable since it either leads to (i) the quote being rejected if this is the customer's first contact, yielding zero payoff (the same outcome as when she quotes $1 + \Delta$), or (ii) the customer accepting with probability one if this is the second quote she receives (a strictly better outcome than when the dealer quotes $1 + \Delta$). \square

We now apply Lemma 2 to show Proposition 1. Since dealers only get contacted at most once in equilibrium, we drop the n th superscript and refer to their on-path quotes as $b_\sigma \equiv b_\sigma^1$. We start by showing that, when $s > 1 - q$, there cannot be search in equilibrium. Otherwise, by Lemma 2, it must be that the customer is willing to search after observing a quote of $b_h = 1 + \Delta$, which gives the condition

$$1 + \Delta - b_h \leq 1 + \Delta - qb_h - (1 - q)b_u - s \iff s \leq 1 - q,$$

a contradiction.

Once we establish that the customer does not search in equilibrium, the ensuing game is analogous to the static game described in Section 3.1. In particular, the absence of search implies that dealers do not learn anything about each other's type, which makes the game starting in every period a repetition of the static one. Letting $\lambda_u(Z)$ and $\lambda_h(Z)$ denote the equilibrium probabilities that the uninformed and the high-type dealers trade conditional on $Z_t = Z$, it follows from (IC $_{l \rightarrow u}$) and (IC $_{u \rightarrow h}$) that $\lambda_h(Z) = \lambda_u(Z) = 0$ for all $Z \in \{0, 1\}$ in the least efficient equilibrium, and $\lambda_u(Z) = \frac{\Delta}{1+\Delta}$ and $\lambda_h(Z) = \left(\frac{\Delta}{1+\Delta}\right)^2$ for all $Z \in \{0, 1\}$ in the most efficient one.

Finally, we argue that satisfying the incentive constraints (IC $_{l \rightarrow u}$), (IC $_{l \rightarrow h}$), (IC $_{u \rightarrow h}$) is sufficient to construct an equilibrium. We can specify off-path beliefs for the customer by making them as low as possible given the belief monotonicity refinement (e.g., the customer believes that the dealer is uninformed after observing any quote in $(\Delta, 1 + \Delta)$). Then, it is easy to check that no dealer type has a profitable deviation, and that each customer is making optimal decisions given her belief.

Proof of Proposition 2

Step 1: Construction of the least efficient equilibrium. First, we construct a no-search equilibrium in which the uninformed and high types never trade. To do this, suppose that, in every period, the uninformed dealer quotes $b_u \in (\max\{1 + \Delta - s, \Delta\}, 1 + \Delta)$ and the high-type dealer quotes $b_h = 1 + \Delta$. After observing a quote equal to b_u or b_h , the customer exits the market with probability one. Off-path, customers believe that a quote $b < b_u$ comes from

a low-type dealer with probability one, and a quote $b \in (b_u, b_h)$ comes from an uninformed dealer with probability one.

To show that this is an equilibrium, let us start by checking sequential rationality for the customer. If the first quote that the customer receives is b_u , the payoffs from accepting are $\Delta - b_u < 0$, while the payoffs from searching are $-s$ since both types of informed dealer are quoting $v + \Delta$. Hence, it is optimal for the customer to exit after being quoted b_u . Similarly, after she observes a quote of b_h , the payoffs from accepting the quote are $1 + \Delta - b_h = 0$, while the payoffs from searching are $(1 - q)(1 + \Delta - b_u) - s < 0$, so it is optimal for her to exit the market. Dealers do not benefit from deviating, since any quote in $(-1 + \Delta, 1 + \Delta]$ is rejected by the customer with probability one. Note that, because the low-type dealer trades with probability one in any equilibrium (see Lemma 2), the fact that the other two types trade with zero probability makes this the least efficient equilibrium in the opaque market.

Step 2: Uniqueness of a no-search equilibrium. We show that, when the search cost s is low enough, any equilibrium in which the customer does not search must be the inefficient one described in the previous step. To see this, suppose that the uninformed type trades with positive probability in a no-search equilibrium. This implies that her equilibrium quote is $b_u \leq \Delta$. By Lemma 2 (which we can apply separately to each period), $b_h \geq 1 + \Delta$, which implies that the customer's expected payoff from searching after receiving a quote of b_h is at least $(1 - q)(1 + \Delta - \Delta) - s$. This would contradict the no-search hypothesis if $s < 1 - q$. So it must be that the uninformed type does not trade in any no-search equilibrium. Additionally, it cannot be that the high type trades with positive probability, since then the uninformed type would find a profitable deviation in mimicking the high type. It follows that, when s is small enough, the most efficient equilibrium must either coincide with the no-search equilibrium from step 1 (in which case the equilibrium outcome is unique), or feature search on the equilibrium path.

Step 3: Construction of the most efficient equilibrium. We argue that, provided that search costs are low enough, there exists an equilibrium in which the customer searches on the equilibrium path if and only if $q \leq \frac{2\Delta^2}{1+\Delta}$; we also characterize the most efficient equilibrium within this class.

Suppose first that the customer searches in period zero. Observe that, in the most efficient equilibrium, players must coordinate on the most efficient static equilibrium in period one. This is because dealers do not observe each other's actions, so there is no scope for punishing deviations by choosing a "bad" continuation outcome. To the extent that the selection of continuation equilibria does not affect dealers' dynamic incentives, coordinating on an equilibrium that is inefficient in the static sense is strictly (welfare-) suboptimal.

First, we derive the customer's strategy in period zero on the equilibrium path. We know

from Lemma 2 that, because there is search in equilibrium, $b_u = \Delta$ and $b_h = 1 + \Delta$ in period zero. The customer's payoff from searching after receiving a quote of b_h is $1 - q - s$, which is positive whenever s is small, so the customer searches with probability one after this quote. After observing a first quote equal to b_u , the customer must exit the market with positive probability, in order to deter the low-type from deviating to Δ . After observing a sequence of quotes (b_h, b_u) , the customer's payoff from trading at b_u is strictly positive and hence she will accept this second quote with certainty. After observing a pair of quotes equal to (b_h, b_h) , by part (v) of Lemma 2, the customer accepts the second quote with probability one.

Second, we show that the outcome in period one must be the one with no search in which only the low-type dealer trades. By the argument in the previous paragraph, we have effectively four types of dealers at the start of period one: the informed low and high types, and two "types" of uninformed dealer. The first type of uninformed dealer is a dealer who either was not contacted or was contacted and did not trade in period zero. This dealer type believes that the value of the asset is zero, since the event of not being contacted or having a quote be rejected by the customer is equally likely under both asset values (considering that, in the event of not being contacted, the dealer believes the customer to be a buyer or a seller with equal probabilities). The second type is given by the dealer who was contacted and traded in period zero, whose belief about the value of the asset will be $q/(2\lambda_u + q)$, by Bayes' rule. Let us index these two types by u_0 and u_{ct} respectively, and let b_σ be an equilibrium quote of type σ .³⁹

Suppose by contradiction that the customer searches with positive probability in period one. Then, it must be that $b_{u_0} < b_{u_{ct}} < b_h \leq 1 + \Delta$. When the search cost is low, this implies that the customer will search with probability one after observing a quote equal to b_h (hoping to find an uninformed type), and after $b_{u_{ct}}$ (hoping to find a u_0 type). There is no search following b_{u_0} . As a result, the u_{ct} type can only trade if she is visited second. Since there can only be one type- u_{ct} dealer in the market, this means that she only trades conditional on the customer having visited a high-type dealer before. But if this is the case, the u_{ct} -dealer can deviate to quoting $b_h > b_u$ which the customer would want to accept for sure. Hence, there cannot be search in equilibrium in period one.

Third, we describe dealers' incentive constraints in period zero. Let λ_u be the probability that the customer agrees to trade after receiving a first quote equal to Δ . Note that, for all dealer types, continuation payoffs are constant in the quote that they provide in period zero. Hence, the condition ensuring that the uninformed dealer does not benefit from deviating to

³⁹Note that dealer's beliefs do not satisfy the martingale condition, that is, they do not average out to the prior. This is because we assume that the customer is a buyer. In the symmetric case of the customer being a seller, dealer's expected value of the asset would be zero or negative in the above analysis.

$1 + \Delta$ is

$$\lambda_u \Delta + \frac{q}{2}(\Delta - 1) \geq \frac{q}{2}\Delta \iff \lambda_u \geq \frac{q}{2\Delta},$$

and the condition ensuring that the low-type dealer does not deviate to Δ is

$$\Delta \geq \lambda_u(1 + \Delta) \iff \lambda_u \leq \frac{\Delta}{1 + \Delta}.$$

Therefore, a necessary condition for an equilibrium with search in period zero to exist is that $\frac{q}{2\Delta} \leq \frac{\Delta}{1 + \Delta} \iff q \leq \frac{2\Delta^2}{1 + \Delta}$, which is the condition we give in Proposition 2. To guarantee that the above conditions are also sufficient, we set off-path beliefs so that the customer believes that the value of the asset is -1 after any first quote below Δ , and that the value is 0 after any first quote between Δ and $1 + \Delta$, which guarantees that all off-path quotes above $-1 + \Delta$ are rejected by the customer with probability one.

The most efficient out of all the equilibria with search in period zero is the one where $\lambda_u = \frac{\Delta}{1 + \Delta}$. Because the welfare objective assigns more weight to period zero than to period one, this equilibrium is more efficient than the alternative search equilibrium in which the search equilibrium is played in period one, and there is no search in period zero. It is also more efficient than the unique (outcome in the) no-search equilibrium and hence is overall the most efficient equilibrium in this regime.

If $q > \frac{2\Delta^2}{1 + \Delta}$, there cannot be search in either period in equilibrium. Therefore, in this case, the most efficient equilibrium coincides with the least efficient one in which only the low-type trades.

Proof of Proposition 3

First, both regimes share the same least efficient equilibrium. It is one where, in every period, the low-type dealer trades with probability one (conditional on being contacted) and the other two types trade with probability zero. Hence,

$$\underline{\Lambda}_{\text{opaq}_{\text{low}}} = \underline{\Lambda}_{\text{opaq}_{\text{high}}} = \frac{q}{2}.$$

To find the most efficient equilibrium, we can apply the bounds derived in the proof of Proposition 1 to write:

$$\bar{\Lambda}_{\text{opaq}_{\text{high}}} = \frac{q}{2} + (1 - q)\frac{\Delta}{1 + \Delta} + \frac{q}{2}\left(\frac{\Delta}{1 + \Delta}\right)^2.$$

Also, as shown in the proof of Proposition 2, if $q > \frac{2\Delta^2}{1+\Delta}$, there only exist no-search equilibria in the opaque market with low search cost. Hence, in that case, $\bar{\Lambda}_{\text{opaqlow}} = \underline{\Lambda}_{\text{opaqlow}} < \bar{\Lambda}_{\text{opaqlow}}$. So we obtain that the high-search-cost opaque market strongly welfare-dominates the low-search one if $q > \bar{q} \equiv \frac{2\Delta^2}{1+\Delta}$.

Consider the case $q \leq \bar{q}$. We showed that the most efficient equilibrium in the low-search-cost scenario satisfies

$$\bar{\Lambda}_{\text{opaqlow}} = \alpha \left(q + (1-q) \frac{\Delta}{1+\Delta} \right) + (1-\alpha) \frac{q}{2},$$

where α was defined in (B.1). Taking the difference,

$$\bar{\Lambda}_{\text{opaqlow}} - \bar{\Lambda}_{\text{opaqlow}} < 0 \iff q < \underline{q} \equiv \frac{2\Delta(1-\alpha)(1+\Delta)}{2\Delta + \alpha + (1-\alpha)\Delta^2},$$

One can also check that Assumption (3.1) implies that $\underline{q} < \bar{q}$. This establishes weak welfare dominance of the high-search-cost regime when $q < \underline{q}$, and the low-search-cost regime when $q \in (\underline{q}, \bar{q}]$.

Proof of Proposition 4

We proceed backwards by first deriving continuation equilibria in period one, then using this result to characterize equilibrium strategies in period zero, and finally showing that they conform to the static monopoly trading pattern.

Step 1: Continuation equilibrium in period one. There are three kinds of on-path period-one histories to consider determined by the transaction or lack thereof announced at the end of period zero. First, if an informed dealer (either with high or low type) transacted in period zero, then, because equilibrium quotes are fully revealing, the market becomes symmetrically informed about the value of the asset. Under the never-dissuaded-once-convinced refinement, the game that ensues is one of sequential search with complete information as in Diamond (1971) in which the unique equilibrium features trade with probability one at the monopolistic price. In this equilibrium, both dealers provide the same quote, and therefore the customer is indifferent between trading with either one of them. Since she may now observe the identity of the dealer who traded in period zero, she may condition the probability that she contacts dealer i first on whether or not dealer i traded in period zero. We denote by ρ^l and ρ^h the probability that the customer first contacts the dealer who traded in period zero, conditional on observing a transaction at a price equal to, respectively, the low- and the high-type's equilibrium price.

Second, let us look at the continuation game after an uninformed-type dealer, say dealer

A , transacted in period zero. Conditional on there not being search in equilibrium in period zero (we can always set the search cost to be high enough to ensure this), dealer A learns nothing in period zero about the value of the asset. Dealer B learns that her opponent is uninformed with probability one, but her estimate of the value of the asset remains unchanged. The period-one customer enters the market with asymmetric beliefs about dealers' types: she knows that dealer A is uninformed, but her belief about the type of dealer B is equal to the prior. Again, if s is high enough, it is without loss of generality to look at equilibria in which customers do not search in period one. This implies that, conditional on being contacted, dealer B with belief $v' \in \{-1, 0, 1\}$ is going to quote $v' + \Delta$, and the low-type dealer B trades with probability one. To ensure that the uninformed type does not want to mimic the high type and that the low type does not want to mimic the uninformed type, we need precisely the same conditions as we derived for the opaque high-search-cost market. That is, letting $\lambda_u(b_u), \lambda_h(b_u)$ be, respectively, the probabilities that the quote of the uninformed and the high types is accepted by the customer at this node of the game, we must have

$$\lambda_u(b_u) \leq \frac{\Delta}{1 + \Delta} \quad \text{and} \quad \lambda_h(b_u) \leq \lambda_u(b_u) \frac{\Delta}{1 + \Delta}.$$

As before, these conditions are also sufficient for incentive compatibility of dealer B if we set off-path beliefs appropriately. As for dealer A , we argue that she must quote Δ and this quote must be accepted by the customer with probability one. This is because dealer A 's type is known by the customer and, since the customer does not search in equilibrium, she will accept any quote from dealer A as long as it is weakly below her willingness to pay Δ .

Next, let us verify conditions on the search cost such that not searching is indeed optimal for the customer at this node of the game. Given dealers' quoting strategies, searching may only be worthwhile after having contacted high-type dealer B , in which case the customer's expected payoff from searching is $1 - s$. So searching is suboptimal whenever $s > 1$. Moreover, the customer is indifferent between contacting dealer A or dealer B . So we have an additional equilibrium object which is the probability that dealer A is contacted first by the customer at this node of the game, which we denote by ρ^u .

Third, let us look at the continuation equilibrium following no trade in period zero. Because of symmetry (across asset values and desired positions of the customer), this event is equally likely to happen when the asset value is either high or low, so it is uninformative about the value of the asset and about the type of the dealer who was contacted. Hence, dealers enter this part of the game with the same belief that they had in period zero, and the period-one customer's belief is equal to the prior. The game is then the same as in

the static opaque market, and the set of equilibrium outcomes are described by the static monopoly trading pattern. In particular, the uninformed and the high-type dealers trade with probabilities $\lambda_u(nt)$, $\lambda_h(nt)$, satisfying $\lambda_u(nt) \leq \frac{\Delta}{1+\Delta}$ and $\lambda_h(nt) \leq \frac{\Delta}{1+\Delta} \lambda_u(nt)$.

Step 2: Equilibrium strategies in period zero. We will argue that dealers quoting strategies must coincide with their static strategy from Lemma 2. First, the same arguments as in the proof of Lemma 2 imply that the low type must quote $-1 + \Delta$ and trade with probability one. Second, suppose that the high type provides a quote $b_h < 1 + \Delta$ in equilibrium. This quote must be accepted with probability one by the customer and lead to a continuation payoff of $\rho^h \Delta$. Consider now her payoff from deviating to $\tilde{b} = b_h + \epsilon < 1 + \Delta$, where $\epsilon > 0$ is small. By monotonicity of beliefs and stationarity of equilibrium strategies, this quote is accepted with probability one by the current customer and leads to the same continuation payoff as b_h (since it leads to the same hierarchy of beliefs), a contradiction. Hence, $b_h = 1 + \Delta$.

Third, we argue that, if the uninformed dealer quotes $b_u < \Delta$, the incentive-compatibility constraint of either the low or the uninformed type must be violated. To see this, note that the low type dealer's equilibrium payoff is $\Delta(1 + \delta\gamma\rho^l)$. If she deviates to b_u , she gets to trade with probability one in the current period and ensures a continuation payoff of $\delta\gamma\rho^u(1 + \Delta)$. Hence, her IC constraint requires that

$$\Delta(1 + \delta\gamma\rho^l) \geq b_u + 1 + \delta\gamma\rho^u(1 + \Delta).$$

Now, consider the uninformed dealer's payoff from deviating to $\tilde{b}_u = \Delta - \epsilon > b_u$, where $\epsilon > 0$ is arbitrarily small. By monotonicity of beliefs, this quote is accepted with certainty by the current customer and leads to a non-negative continuation payoff. In equilibrium, on the other hand, the uninformed dealer is trading with probability one in the current period at a price equal to b_u , and gets a continuation payoff equal to $\delta\gamma\rho^u \Delta$. Hence, a necessary condition for the uninformed dealers' incentive compatibility is

$$b_u + \delta\gamma\rho^u \Delta \geq \Delta.$$

For the above two conditions to be jointly satisfied, we need $\delta\gamma\rho^l \Delta - 1 - \delta\gamma\rho^u \geq 0$, which is a contradiction with assumption (3.1). Therefore, the uninformed type quotes Δ with probability one.⁴⁰

Given these quoting strategies, let λ_u and λ_h denote the probability that the customer accepts, respectively, a quote equal to Δ and $1 + \Delta$ in period zero. By looking at the

⁴⁰As before, we do not need to explicitly consider the case in which either the high or the uninformed types quote a price strictly above the buyer's valuation. Such an equilibrium is outcome-equivalent to one where she follows the strategy described above and the customer always rejects.

incentive-compatibility constraints of the dealers, we can again derive the relevant bounds on equilibrium values of λ_u and λ_h . Let us start by checking the condition ensuring that the uninformed dealer does not benefit from deviating to $1 + \Delta$. If the uninformed dealer deviates to $1 + \Delta$, her optimal strategy in period one is as follows. If the quote is accepted, then she can trade with probability one at a quote of $1 + \Delta$, so this will be her optimal continuation strategy. If the quote is rejected, our construction of equilibrium in the continuation game following no trade ensures that it is optimal for the uninformed dealer to quote Δ and trade with probability $\lambda_u(nt)$. Hence, her incentive-compatibility constraint can be written as

$$\lambda_u \Delta + \frac{\delta \gamma}{2} (\lambda_u 2 \rho^u + (1 - \lambda_u) \lambda_u(nt)) \Delta \geq \lambda_h (1 + \Delta) + \frac{\delta \gamma}{2} (\lambda_h 2 \rho^h (1 + \Delta) + (1 - \lambda_h) \lambda_u(nt) \Delta).$$

It remains to check the incentive constraint of the low-type dealer. The only relevant one is that she does not benefit from quoting Δ (in particular, the IC constraint of the uninformed type ensures that the low type does not want to mimic the high type). If she deviates to Δ and her quote is accepted in period zero, she can trade with probability ρ^u (the probability of being contacted) at the same price in period one. Otherwise, if the quote is rejected, her optimal continuation strategy is to quote $-1 + \Delta$. The resulting incentive constraint is

$$\Delta + \delta \gamma \rho^l \Delta \geq \lambda_u (1 + \Delta) + \frac{\delta \gamma}{2} (2 \rho^u \lambda_u (1 + \Delta) + (1 - \lambda_u) \Delta).$$

Finally, given dealers' quoting strategies, not searching is optimal for the customer whenever $s > 1 - q$. This establishes that the static monopoly trading pattern must emerge in period zero.

The above conditions are also sufficient for an equilibrium if we set off-path beliefs in the usual way.

Proof of Proposition 5

Throughout, we use the already established fact that, in equilibrium, conditional on being contacted, the low type must provide a quote equal to $-1 + \Delta$, and this quote must be accepted by the customer with probability one.

Step 1: Construction of the least efficient equilibrium. We set the uninformed and the high type's quoting strategies in period zero to be as in the least efficient equilibrium from Proposition 2, which ensures that there is no search and only the low type trades in period zero. In period one, there are two possible on-path continuation games depending on the outcome in period zero: either there was no trade, or there was trade at a price equal to $b_l = -1 + \Delta$. In

the former case, since the event of no trade is uninformative about the value of the asset and about dealers' types, the game that ensues is a one-shot version of the static search game described in Lemma 2, for which we know that there always exists a no-search equilibrium in which only the low type trades. In the latter case, the market becomes symmetrically informed about the asset's value, and the unique equilibrium is the Diamond-type equilibrium where there is no search and trade happens with probability one at a price equal to $-1 + \Delta$. By our arguments in the proof of Proposition 2, this is the unique equilibrium featuring no search in period zero.

Step 2: Construction of the most efficient equilibrium. We proceed backwards and construct an equilibrium where the customer searches with positive probability in period zero. We show that this is the most efficient equilibrium and provide conditions for its existence.

2.a) Continuation equilibrium in period one. As in Proposition 4, following trade at a price that fully reveals the state, the only possible outcome in period one is efficient and features no search.

If the outcome in period zero was trade at a price coming from the uninformed type that does not fully reveal the state,⁴¹ we argue that the only continuation equilibrium is such that: (i) the dealer who traded in period zero gets contacted by the customer with zero probability, and (ii) the static monopoly trading pattern is played (with the dealer who did not trade in period zero acting as the monopolist). To show this, suppose that the uninformed dealer A traded in period zero at a non-fully-revealing quote. In period one, the customer knows with certainty that dealer A has the uninformed type. Moreover, the type of dealer B is unknown and hence, if contacted, dealer B will quote a price equal to her belief about the asset's value plus Δ .

Suppose first that dealer A provides a quote $b_u \leq \Delta$ with positive probability in this part of the game. Then, assuming low enough search cost, the optimal search strategy for the customer would be to always contact dealer B first, and to request a second quote from dealer A only when dealer B quoted $1 + \Delta$. But, if this is the case, dealer A knows that she gets contacted only when the state is high, and her optimal strategy is to quote $1 + \Delta$, contradicting $b_u \leq \Delta$. So it must be that dealer A provides only quotes above Δ in equilibrium. Again, since this quote only leads to trade after the customer contacts a high type first, the optimal quote for dealer A will be $1 + \Delta$. Under this quoting strategy by the dealers, the optimal search pattern for the customer is to always contact dealer B first and

⁴¹This price might still reveal some partial information about the state. This is because, in principle, the uninformed dealer may get contacted more than once, and her quoting strategy at different rounds may reflect the information that she gathered from the customer's searching behavior. The case in which the uninformed type eventually becomes fully informed and her quote fully reveals the state is addressed in the previous paragraph. Repeated contacts will be ruled out later in the proof.

never contact dealer A .

Thus, the set of possible continuation equilibrium outcomes after trade at a price equal to the uninformed type's price can be constructed as follows. Conditional on being contacted (which is an off-path event), dealer A believes that the asset's value is 1, and provides a quote equal to $1 + \Delta$. Dealer B provides a quote equal to the monopolistic price. The game that ensues when dealer B is contacted is the same as in the static monopoly benchmark. In particular, by our discussion in Section 3.1, the low-type must trade with probability one, and the uninformed and high types trade with probability $\lambda_u(b_u)$ and $\lambda_h(b_u)$ satisfying

$$\lambda_u(b_u) \leq \frac{\Delta}{1 + \Delta}, \quad \lambda_h(b_u) \leq \frac{\Delta}{1 + \Delta} \lambda_u(b_u).$$

As before, these conditions become sufficient by choosing off-path beliefs that lead the customer to reject any off-path quote above $-1 + \Delta$.

It remains to characterize the continuation outcome following no trade in period zero. We show that this outcome must be analogous to the inefficient one that arises in the fully opaque market in period one. To do so, we must first describe the on-path search behavior in period zero. By very similar arguments to the ones provided in the proof of Lemma 2, which we omit to avoid repetition, we can show that the searching path and dealers' quoting strategies coincide with the static description from Lemma 2.

This implies that in period zero there is search only following a quote from the high type, and thus that the continuation game that ensues after no trade is exactly the same as described in the proof of Proposition 2. In particular, the event of no trade is uninformative for the customers, and hence the buyer in period one will hold a belief equal to the prior. Moreover, there will be two different types of the uninformed dealer holding different beliefs about the asset's value, and the only possible continuation equilibrium is the no-search one where only the low type trades.

2.b) Equilibrium strategies in period zero. First, we show that the quoting strategy of the uninformed type must be to quote Δ with probability one. Suppose by contradiction that she quotes $b_u < \Delta$ with positive probability. Since there is no search following b_u , this quote is accepted with probability one and our preceding analysis implies that her continuation payoff is 0. But then the uninformed dealer has a profitable deviation, which is to slightly increase her quote: by belief monotonicity, this higher quote is still accepted by the customer with probability one, and continuation payoffs can only weakly improve. By a similar argument, the high type must quote $1 + \Delta$ with probability one. If her strategy involves quoting $b_h \in [1, 1 + \Delta)$, this quote gets rejected with probability one by a first-contact customer if s is low enough. This is because the customer expects to find an uninformed dealer quoting

$\Delta < b_h$. Moreover, if this is the customer's second contact, she will accept for sure any quote less than $1 + \Delta$. It follows that the high type can strictly increase her payoff by slightly increasing her quote. This makes her strictly better off against a second-contact customer in period zero while keeping her continuation payoff unchanged, and without hurting her payoff in the event that this is instead the customer's first contact.

Having described quoting strategies, we write down dealers' incentive constraints in period zero. Let λ_u be the probability that a quote coming from the uninformed type is accepted. Consider first the condition ensuring that the low type does not benefit from deviating to Δ . Following this deviation, her optimal continuation strategy is as follows: if her quote is accepted in period zero, then she will be contacted by the customer with zero probability and her payoff is 0; if her quote is rejected, she will be contacted with probability $1/2$, and her optimal strategy is to quote $-1 + \Delta$ and trade with probability one. The resulting incentive constraint is thus

$$\Delta + \delta\gamma\rho^l\Delta \geq \lambda_u(1 + \Delta) + \delta\frac{\gamma(1 - \lambda_u)}{2}\Delta.$$

In the most efficient equilibrium, $\rho^l = 1$, and thus $\lambda_u = \frac{\Delta(1+\delta\gamma/2)}{1+(1-\delta\gamma/2)\Delta} \equiv \bar{\lambda}_u(\text{post}_{\text{low}})$.

Next, let us derive a condition preventing an informed dealer, say dealer A , from deviating to $1 + \Delta$. Consider first her optimal continuation payoff following this deviation. If the outcome was trade at a high quote (which only happens if the opponent has a high type), she will quote $1 + \Delta$ and trade with probability ρ^h . If there is no trade, then the continuation equilibrium is the no-search one where the uninformed dealer gets zero payoff. If the outcome is trade at the uninformed type's price (i.e., dealer B is uninformed and gets contacted second), then dealer A will be contacted in the following period with probability one and trade at a price equal to Δ with probability $\lambda_u(b_u)$. Finally, there is the possibility that, after receiving a high quote from dealer A and searching, the customer encounters a low-type dealer B who provides a quote equal to $-1 + \Delta$. Since observing a pair of quotes equal to $(1 + \Delta, -1 + \Delta)$ is an off-path event, we must specify the customer's beliefs and her strategy following this history. No matter what the customer's beliefs are, she will weakly prefer to accept the second quote of $-1 + \Delta$, and she will strictly want to do so unless her belief about its value is equal to -1 . Moreover, if the customer accepts the second quote, the uninformed dealer's continuation payoff is $(1 - \rho^l)\Delta$, since she will quote $-1 + \Delta$ in period one. Otherwise, if the outcome is no trade, her continuation payoff will be 0. In order to sustain the largest possible set of equilibria, we consider off-path beliefs which minimize the payoffs from deviating, which in this case requires setting the customer's belief after observing $(1 + \Delta, -1 + \Delta)$ to -1 and her strategy to rejecting the second quote with probability one.

On the other hand, dealer A 's continuation payoff when she quotes her equilibrium strategy Δ and the state is persistent in the following period is 0 by our preceding analysis. Putting everything together, we have that the low type's incentive constraint is

$$\lambda_u \Delta + \frac{q}{2}(\Delta - 1) \geq \frac{q}{2}\Delta + \frac{\delta\gamma}{2} \left(\frac{q}{2} + (1-q)\lambda_u(b_u) \right) \Delta.$$

This yields the following necessary condition for existence of an equilibrium with search:

$$\begin{aligned} \frac{\Delta(1 + \frac{\delta\gamma}{2})}{1 + (1 - \frac{\delta\gamma}{2})\Delta} &\geq \frac{\frac{q}{2}(1 + \frac{\delta\gamma}{2}\Delta) + \frac{\delta\gamma}{2}(1-q)\lambda_u(b_u)\Delta}{\Delta}, \quad \text{for some } \lambda_u(b_u) \in \left[0, \frac{\Delta}{1 + \Delta}\right] \\ &\iff q \leq \frac{2\Delta^2(1 + \frac{\delta\gamma}{2})}{(1 + (1 - \frac{\delta\gamma}{2})\Delta)(1 + \frac{\delta\gamma}{2}\Delta)}. \end{aligned}$$

Finally, note that, for every dealer type, the continuation payoff if there is no trade in period zero is weakly less than the continuation payoff conditional on trading. This implies that all quotes that lead to no trade in period zero are not a profitable deviation relative to the quote they are playing in equilibrium. Hence, the above conditions are also sufficient for an equilibrium if we set customer's off-path beliefs in the usual way so that she refuses to trade at any off-path quote above $-1 + \Delta$.

This concludes the construction of the most efficient equilibrium among those that feature search in period zero. By the same arguments as in Proposition 2, this is the most overall efficient equilibrium, provided that it exists. When $q > \frac{2\Delta^2(1+\delta\gamma/2)}{(1+(1-\delta\gamma/2)\Delta)(1+\delta\gamma\Delta/2)}$, a search equilibrium does not exist, and the unique equilibrium outcome is the least efficient no-search equilibrium described above.

Proof of Proposition 6

We begin with the market with low search cost. Applying the equilibrium characterization from Proposition 5, we can write the welfare lower bound as

$$\underline{\Lambda}_{\text{post}_{\text{low}}} = \alpha \frac{q}{2} + (1 - \alpha) \left(1 + \frac{1}{2}(1 - q) \right) \frac{q}{2},$$

where α was defined in (B.1). Thus, $\underline{\Lambda}_{\text{opaq}_{\text{low}}} = \frac{q}{2} < \underline{\Lambda}_{\text{post}_{\text{low}}}$.

Next, consider the most efficient equilibrium in the post-trade transparent regime. Let $\bar{q}_{\text{opaq}_{\text{low}}} \equiv \frac{2\Delta^2}{1+\Delta}$ and $\bar{q}_{\text{post}_{\text{low}}} \equiv \frac{2\Delta^2(1+\delta\gamma/2)}{(1+(1-\delta\gamma/2)\Delta)(1+\delta\gamma\Delta/2)}$ be the cutoffs on q below which an equilibrium with search exists in each of the two regimes. From the proof of Proposition 5, ff $q \leq \bar{q}_{\text{post}_{\text{low}}}$, then the most efficient equilibrium features search in period zero and yields

welfare

$$\begin{aligned} \bar{\Lambda}_{\text{post}_{\text{low}}} &= \alpha(q + (1 - q)\bar{\lambda}_u(\text{post}_{\text{low}})) + (1 - \alpha)\left(\frac{q}{2}(1 + q) + \frac{q}{2}(1 - q)\lambda_h(b_u)\right. \\ &\quad \left. + (1 - q)\bar{\lambda}_u(\text{post}_{\text{low}})\left(\frac{q}{2}(1 + \lambda_h(b_u)) + (1 - q)\lambda_u(b_u)\right) + (1 - q)(1 - \bar{\lambda}_u(\text{post}_{\text{low}}))\frac{q}{2}\right), \end{aligned}$$

with $\lambda_u(b_u) = \Delta/(1 + \Delta)$ and $\lambda_h(b_u) = (\Delta/(1 + \Delta))^2$. To compare welfare across the two regimes, observe that, if $q > \bar{q}_{\text{post}_{\text{low}}} > \bar{q}_{\text{opaqlow}}$, then an equilibrium with search does not exist in either regime, so in that case: $\underline{\Lambda}_{\text{opaqlow}} = \bar{\Lambda}_{\text{opaqlow}} < \underline{\Lambda}_{\text{post}_{\text{low}}} = \bar{\Lambda}_{\text{post}_{\text{low}}}$, and hence the post-trade transparent market strongly welfare-dominates the opaque one.

On the other hand, if $q \leq \bar{q}_{\text{opaqlow}}$, then both regimes have an equilibrium with search. We can verify that the post-trade transparent market dominates the opaque one, in terms of its most efficient equilibrium, in every period (i.e., for every realization of $Z \in \{0, 1\}$). To do so, let $\bar{\Lambda}_{\text{post}_{\text{low}}}(Z)$ and $\bar{\Lambda}_{\text{opaqlow}}(Z)$ be the per-period probabilities of trade under the most efficient equilibrium in each of the two regimes. Using our characterization, we can write

$$\begin{aligned} \bar{\Lambda}_{\text{post}_{\text{low}}}(0) - \bar{\Lambda}_{\text{opaqlow}}(0) &= (1 - q)\left(\frac{\Delta(1 + \frac{\delta\gamma}{2})}{1 + (1 - \frac{\delta\gamma}{2})\Delta} - \frac{\Delta}{1 + \Delta}\right) > 0, \\ \bar{\Lambda}_{\text{post}_{\text{low}}}(1) - \bar{\Lambda}_{\text{opaqlow}}(1) &= \frac{q^2}{2} + \frac{q}{2}(1 - q)\lambda_h(b_u) + (1 - q)\lambda_u\left(\frac{q}{2}(1 + \lambda_h(b_u))\right. \\ &\quad \left. + (1 - q)\lambda_u(b_u)\right) + (1 - q)(1 - \lambda_u)\frac{q}{2} > 0. \end{aligned}$$

Thus, the post-trade transparent market with low search cost weakly welfare-dominates its opaque counterpart when $q \leq \bar{q}_{\text{opaqlow}}$.

Finally, if $q \in (\bar{q}_{\text{opaqlow}}, \bar{q}_{\text{post}_{\text{low}}}]$, only the post-trade transparent market has a search equilibrium. Therefore, $\bar{\Lambda}_{\text{post}_{\text{low}}} > \underline{\Lambda}_{\text{post}_{\text{low}}} > \underline{\Lambda}_{\text{opaqlow}} = \bar{\Lambda}_{\text{opaqlow}}$. Therefore, strong welfare dominance obtains when $q > \bar{q}_{\text{opaqlow}}$.

We now turn to the market with high search costs. The least efficient equilibrium in the post-trade transparent market has the same trading outcome as in the low-search-cost setting. Namely, only the low type trades in every on-path history, except for period one if trade happened in period zero, in which case the outcome is efficient. Hence,

$$\underline{\Lambda}_{\text{post}_{\text{high}}} = \alpha\frac{q}{2} + (1 - \alpha)\left(1 + \frac{1}{2}(1 - q)\right)\frac{q}{2} > \frac{q}{2} = \underline{\Lambda}_{\text{opaqlow}_{\text{high}}}.$$

Next, we consider the upper bound on welfare in the post-trade transparent regime. By

the proof of Proposition 4, denoting $\lambda = (\lambda_u, \lambda_h, \lambda_u(b_u), \lambda_h(b_u), \lambda_u(nt), \lambda_h(nt), \rho^l, \rho^u, \rho^h)$, the upper bound solves:

$$\begin{aligned}
(1 + \delta\gamma)\bar{\Lambda}_{\text{post high}} &= \max_{\lambda \in [0,1]^9} \left(\frac{q}{2}(1 + \lambda_h) + (1 - q)\lambda_u \right) + \delta\gamma \left\{ \frac{q}{2}(1 + \lambda_h) \right. \\
&\quad + \frac{q}{2}(1 - \lambda_h) \frac{(1 + q)\lambda_h(nt) + (1 - q)\lambda_u(nt)}{2} \\
&\quad + (1 - q)\lambda_u \left[\rho + (1 - \rho) \left(\frac{q}{2}(1 + \lambda_h(b_u)) + (1 - q)\lambda_u(b_u) \right) \right] \\
&\quad \left. + (1 - q)(1 - \lambda_u) \frac{\frac{q}{2}(1 + \lambda_h(nt)) + (2 - q)\lambda_u(nt)}{2} \right\} \quad (\text{Opt}_{\text{post high}}) \\
\text{s.t. } \Delta + \delta\gamma\rho^l\Delta &\geq \lambda_u(1 + \Delta) + \frac{\delta\gamma}{2}(2\rho^u\lambda_u(1 + \Delta) + (1 - \lambda_u)\Delta), \quad (\text{ICl}(0)) \\
\lambda_u\Delta + \frac{\delta\gamma}{2}(\lambda_u 2\rho^u + (1 - \lambda_u)\lambda_u(nt))\Delta &\geq \\
\lambda_h(1 + \Delta) + \frac{\delta\gamma}{2}(\lambda_h 2\rho^h(1 + \Delta) + (1 - \lambda_h)\lambda_u(nt)\Delta), &\quad (\text{ICu}(0)) \\
\lambda_u(j) &\leq \frac{\Delta}{1 + \Delta}, \quad j = b_u, nt, \quad (\text{ICl}(1, j)) \\
\lambda_h(j) &\leq \frac{\Delta\lambda_u(j)}{1 + \Delta}, \quad j = b_u, nt. \quad (\text{ICu}(1, j))
\end{aligned}$$

Without deriving an explicit solution to the above program, we can obtain a lower bound on $\bar{\Lambda}_{\text{post high}}$ which is weakly higher than $\bar{\Lambda}_{\text{opaq high}}$. To do so, let $\tilde{\Lambda}_{\text{post high}}$ be the value of the objective when we set the vector of control variables λ to be such that $\rho^l = \rho^h = \rho^u = \frac{1}{2}$ and all the incentive-compatibility constraints in the program hold with equality. Since this is a feasible value of λ , it follows that $\tilde{\Lambda}_{\text{post high}} \leq \bar{\Lambda}_{\text{post high}}$. Moreover, taking the difference with respect to the opaque market, one can verify that $\tilde{\Lambda}_{\text{post high}} - \bar{\Lambda}_{\text{opaq high}} > 0$, and therefore $\bar{\Lambda}_{\text{post high}} - \bar{\Lambda}_{\text{opaq high}} > 0$.

Proof of Proposition 7

First, we argue that the high-search-cost regime strongly welfare-dominates the low-search-cost one when $q > \bar{q}_{\text{post low}}$. This is because, in this region, the only equilibrium that exists when the search cost is low is the inefficient no-search one where only the low-type dealer trades. As a result, $\bar{\Lambda}_{\text{post low}} = \underline{\Lambda}_{\text{post low}} = \underline{\Lambda}_{\text{post high}} < \bar{\Lambda}_{\text{post high}}$.

Next, we show that the high-search-cost regime is weakly welfare-dominant in the limit as $q \rightarrow 0$. Note that, when q is small, an equilibrium with search exists in the market with

low search cost and hence we can write, using our previous results,

$$\lim_{q \rightarrow 0} \bar{\Lambda}_{\text{post low}} = \alpha \cdot \frac{\Delta(1 + \frac{\delta\gamma}{2})}{1 + (1 - \frac{\delta\gamma}{2})\Delta} + (1 - \alpha) \cdot \frac{\Delta(1 + \frac{\delta\gamma}{2})}{1 + (1 - \frac{\delta\gamma}{2})\Delta} \cdot \frac{\Delta}{1 + \Delta}.$$

On the other hand, given our equilibrium conditions in the post-trade transparent market with high search cost, a feasible level of equilibrium welfare obtains from setting the control variables in $(\text{Opt}_{\text{post high}})$ as follows: $\rho^u = 0$, $\rho^l = \rho^h = 1$, $\lambda_u = \frac{\Delta(1+\delta\gamma/2)}{1+(1-\delta\gamma/2)\Delta}$, $\lambda_u(b_u) = \lambda_u(nt) = \frac{\Delta}{1+\Delta}$. Denoting by $\tilde{\Lambda}_{\text{post high}}$ the level of welfare under this equilibrium and taking the limit as $q \rightarrow 0$ we obtain

$$\lim_{q \rightarrow 0} \bar{\Lambda}_{\text{post high}} \geq \lim_{q \rightarrow 0} \tilde{\Lambda}_{\text{post high}} = \alpha \cdot \frac{\Delta(1 + \frac{\delta\gamma}{2})}{1 + (1 - \frac{\delta\gamma}{2})\Delta} + (1 - \alpha) \cdot \frac{\Delta}{1 + \Delta}.$$

Thus,

$$\lim_{q \rightarrow 0} (\bar{\Lambda}_{\text{post high}} - \bar{\Lambda}_{\text{post low}}) \geq (1 - \alpha) \cdot \left(1 - \frac{\Delta}{1 + (1 - \frac{\delta\gamma}{2})\Delta} \right) \cdot \frac{\Delta}{1 + \Delta} > 0.$$

This, together with the fact that $\underline{\Lambda}_{\text{post low}} = \underline{\Lambda}_{\text{post high}}$, establishes the result.

Proof of Proposition 8

Step 1: Continuation equilibrium in period one. Consider the game that ensues in period one following any pair of on-path quotes in period zero. Because under the pre-trade transparent regime dealers observe each other's quotes at the end of each period, both dealers will be symmetrically informed and hold a common belief about the asset's value at the time of quoting. The period-one customer, on the other hand, enters the market with beliefs about v equal to the prior.

Denoting dealers' common belief at this part of the game by $v' \in \{-1, 0, 1\}$, we argue that the unique equilibrium features both dealers quoting v' and trade happening with probability one. To see this, let $\bar{b}(v')$ be the supremum of the support of dealers' quoting strategies. If $\bar{b}(-1) > -1 + \Delta$, then quoting \bar{b} leads to a zero payoff, which is a contradiction since the low-type dealer can secure a strictly positive payoff by quoting a price slightly below $-1 + \Delta$ which is accepted with positive probability. But this implies that it also cannot be that $\bar{b}(0) > \Delta$; otherwise, the uninformed type would be making a zero payoff and could do strictly better by quoting some \tilde{b} slightly below Δ . This is because, since $\bar{b}(-1) < \Delta$, \tilde{b} is an off-path quote, which implies (by our refinement on beliefs) that the customer's belief about v after observing the opponent's quote will be equal to 0 and hence she will be willing to

accept the quote \tilde{b} . An analogous argument shows that $\bar{b}(1) \leq 1 + \Delta$.

If $v' < \bar{b}(v')$ and dealers' strategies assign strictly positive probability to $\bar{b}(v')$, then dealer i can do strictly better by quoting some \tilde{b} slightly below $\bar{b}(v')$, which discontinuously increases dealer i 's probability of trading at a strictly profitable quote. This is because i 's opponent's equilibrium quote will lead the customer to believe that the asset's value is v' , and the customer will be willing to accept i 's quote with probability one whenever the other dealer quotes $\bar{b}(v')$. So dealers' strategies cannot have an atom at $\bar{b}(v')$. But then, quoting approximately $\bar{b}(v')$ leads dealer i to trade with zero probability, which is again a contradiction since she can strictly gain by providing a lower quote. So $\bar{b}(v') \leq v'$. Since dealers would make losses from quoting below v' , it must be that they quote exactly v' with probability one. Given this quoting strategy by the dealers, the customer strictly prefers to trade after being quoted v' , and hence the equilibrium outcome is fully efficient.

Step 2: Equilibrium strategies in period zero. We begin by deriving necessary conditions on dealers' period-zero quoting strategies. We then show that these conditions imply that there are two kinds of equilibria that may arise in period zero: a fully efficient one, and one where the uninformed type trades with interior probability. The latter equilibrium always exists under assumption (3.1), whereas the former one exists if and only if $\delta\gamma \leq \bar{\delta}_{\text{pre}} < \bar{\delta}_{\text{exist}}$.

2.a) Necessary conditions for equilibrium strategies. Let b_σ and \bar{b}_σ be the infimum and supremum of the quoting strategy of a dealer with type $\sigma \in \{l, u, h\}$. Observe that, regardless of her quote in period zero, the continuation payoff of a high-type dealer is 0. This is because, in period one, the opponent will always quote her posterior belief about v which is weakly less than one. Given the lack of effect of today's strategy on future payoffs, we can apply an argument almost identical to the one in the previous paragraph to show that $b_h = \bar{b}_h = 1$.

For the low and the uninformed types, the continuation payoff if they play their equilibrium strategy in period zero is also 0. Since this is their worst-case payoff, it follows that any deviation would lead to a weak improvement in terms of continuation payoffs. Hence, a necessary condition for these two types to not deviate is that flow payoffs in period zero are maximized by any quote in the support of their equilibrium strategy.

Let us now derive the quoting strategy of the low type. We start by noting that $b_l > -1$. Otherwise, the low type makes at most a zero payoff, which is a contradiction since she can do strictly better by quoting a price slightly below $-1 + \Delta$ and trading with probability at least $1 - q$ (since $b_u \geq 0 > -1 + \Delta$ and the customer accepts the lowest quote with probability one whenever that quote is below $-1 + \Delta$). This implies that the low type's equilibrium flow payoff in period zero is strictly positive, and that her equilibrium strategy must be atomless (otherwise, undercutting the price that occurs with strictly positive probability would be a profitable deviation). Also, $\bar{b}_l \leq -1 + \Delta$; otherwise, the low type would get a

zero payoff by quoting \bar{b}_l (which is rejected by the customer for sure). Moreover, it must be that $\bar{b}_l \geq -1 + \Delta$. This is because, if $\bar{b}_l < -1 + \Delta$, the low type's payoff from quoting a price close to \bar{b}_l is approximately $(1 - q)(\bar{b}_l + 1)$. This can be improved upon by quoting a price close to $-1 + \Delta < \underline{b}_u$, which yields a payoff of $(1 - q)\Delta$. Hence, $\bar{b}_l = -1 + \Delta$ and the low type's period-zero payoff is $(1 - q)\Delta$. In order to make her indifferent across all quotes in $[\underline{b}_l, -1 + \Delta]$, we require that

$$(q(1 - F_l(b)) + 1 - q)(b + 1) = (1 - q)\Delta, \quad \forall b \in [\underline{b}_l, -1 + \Delta],$$

where $F_l(b)$ is the cumulative distribution function (cdf) of the low type's quoting strategy. This yields $F_l(b) = 1 - \frac{(1-q)(-1+\Delta-b)}{q(1+b)}$ and $\underline{b}_l = (1 - q)(-1 + \Delta) - q$.

Let us now consider the uninformed dealer. First, suppose that $\bar{b}_u > 1$. Such a quote must be rejected by the customer with probability one; otherwise, the high type dealer, whose equilibrium payoff in period zero is 0, would benefit from deviating to \bar{b}_u . But then the uninformed dealer can do strictly better by quoting 1 which leads to trade at a profit with strictly positive probability (when the opponent is uninformed and quotes a price close to \bar{b}_u). Next, let us show that $\bar{b}_u \leq \Delta$. If $\Delta < \bar{b}_u < 1$, then the uninformed dealer who quotes \bar{b}_u only gets to trade if the opponent's type is high. This is because either the uninformed opponent quotes a price strictly lower than \bar{b}_u (which is preferred by the customer), or she quotes exactly \bar{b}_u (in which case the customer will refuse to trade). But then, the expected payoff from quoting \bar{b}_u is $\frac{q}{2}(\bar{b}_u - 1) < 0$, a contradiction.

In the same way, in order to rule out the uninformed type making a strictly negative payoff, we need that the uninformed type's strategy has an atom of size $\beta > 0$ at \bar{b}_u and that $\bar{b}_u > 0$. We now argue that this implies that $\underline{b}_u = \bar{b}_u$. The uninformed dealer's payoff from quoting \bar{b}_u is $(1 - q)\frac{\beta}{2}\bar{b}_u + \frac{q}{2}(\bar{b}_u - 1)$. Suppose that $\underline{b}_u < \bar{b}_u$ and consider the uninformed dealer's payoff from quoting $b \in (\underline{b}_u, \bar{b}_u)$. Since beliefs are monotone, the customer will believe that this quote comes from the uninformed type, and since $b < \Delta$ the customer will be willing to accept with probability one whenever b is the lowest quote she observes. Hence, the payoff from $b \approx \bar{b}_u$ is approximately $(1 - q)\beta\bar{b}_u + \frac{q}{2}(\bar{b}_u - 1)$ which is strictly higher than her equilibrium payoff. Given this, let $b_u = \underline{b}_u = \bar{b}_u \leq \Delta$.

2.b) Construction of a fully efficient equilibrium. Suppose first that $b_u < \Delta$. Observe that under all pairs of quotes that can arise with positive probability given dealers' quoting strategies, the customer strictly benefits from accepting the lowest quote. Hence, the outcome in such an equilibrium will be fully efficient.

Let us check dealers' incentive constraints in period zero. The high type cannot deviate, since any deviation leads to either no trade in period zero or to trade at a loss. For the

uninformed type, any deviation to a quote less than 1 leads to trade only if the opponent has a high type, and it will lead to a zero payoff in the following period if it is strictly less than 1 since it will lead the opponent to believe that the dealer is uninformed. Hence, the optimal deviation for the uninformed type is to quote 1. After doing so, she can trade at a price equal to 1 with probability 1/2 in the following period. This yields the constraint

$$\frac{1-q}{2}b_u + \frac{q}{2}(b_u - 1) \geq \delta\gamma\frac{1-q}{2} \iff b_u \geq \frac{q + \delta\gamma(1-q)}{1-q} \equiv \underline{b}_u(\text{pre}).$$

The low type's equilibrium payoff in period zero is $(1-q)\Delta$. If she deviates to b_u , she gets $\frac{1-q}{2}(b_u + 1)$ in period zero. If the opponent is uninformed, in period one, the low type will either quote 0 and trade with probability 1/2 or quote $-1 + \Delta$ and trade with probability one (either one of these strategies can be optimal depending on parameter values). As a result, the IC constraint preventing this deviation is

$$(1-q)\Delta \geq (1-q)\frac{b_u + 1}{2} + \delta\gamma\frac{1-q}{2}\kappa \iff b_u \leq 2\Delta - 1 - \delta\gamma\kappa \equiv \bar{b}_u(\text{pre}) < \Delta,$$

where $\kappa = \max\{1, 2\Delta\}$.

For this equilibrium to exist, we need:

$$\underline{b}_u(\text{pre}) \leq \bar{b}_u(\text{pre}) \iff \delta\gamma \leq \frac{(1-q)(2\Delta - 1) - q}{(1-q)(1 + \kappa)} \equiv \bar{\delta}_{\text{pre}} < \bar{\delta}_{\text{exist}}. \quad (\text{B.5})$$

2.c) Construction of an inefficient equilibrium. Finally, let us consider the case with $b_u = \Delta$. After observing a pair of quotes equal to (Δ, Δ) , the customer is indifferent between accepting either quote (with the tie broken uniformly at random) or rejecting both. Let λ_u denote the probability that she accepts one of the quotes. As before, we have to check that the uninformed and low types do not benefit from deviating to, respectively, 1 and b_u . For the uninformed type, we then obtain the following incentive constraint:

$$\frac{(1-q)\lambda_u}{2}\Delta + \frac{q}{2}(\Delta - 1) \geq \delta\gamma(1-q)\frac{1}{2} \iff \lambda_u \geq \frac{\delta\gamma(1-q) + q(1-\Delta)}{(1-q)\Delta} \equiv \underline{\lambda}_u(\text{pre}).$$

On the other hand, the incentive constraint for the low type is

$$(1-q)\Delta \geq \frac{(1-q)\lambda_u}{2}(\Delta + 1) + \frac{\delta\gamma(1-q)}{2}\kappa \iff \lambda_u \leq \frac{2\Delta - \delta\gamma\kappa}{1 + \Delta} \equiv \bar{\lambda}_u(\text{pre}).$$

For this equilibrium to exist, we therefore need

$$\underline{\lambda}_u(\text{pre}) \leq \bar{\lambda}_u(\text{pre}) \iff \delta\gamma \leq \frac{(2-q)\Delta^2 - q}{(1-q)(1+\Delta+\Delta\kappa)}.$$

This condition is implied by our assumption (3.1).

For both kinds of equilibria, the above conditions are also sufficient if we set customer's off-path beliefs in period zero appropriately. To do this, take a one-sided deviation of the form (\tilde{b}, b_σ) , where b_σ is a quote in the support of type σ 's strategy and \tilde{b} is an off-path quote. If $\sigma = l$ or $\sigma = h$, beliefs assign probability one respectively to $v = -1$ or $v = 1$. If $\sigma = u$, we set off-path beliefs as a function of \tilde{b} in the usual way, by setting them equal to -1 if $\tilde{b} < b_u$ and equal to 0 if $\tilde{b} \in (b_u, 1)$. Off-path beliefs conditional on two-sided deviations can be set in any way without affecting the equilibrium conditions.

Proof of Proposition 9

Let us start by showing the first part of the proposition involving weak welfare dominance. Among all equilibria in the pre-trade transparent regime that were identified in Proposition 8, the least efficient one is the one where, in period zero, the uninformed type quotes Δ and the customer agrees to trade after observing a pair of quotes (Δ, Δ) with probability $\underline{\lambda}_u(\text{pre})$. The outcome is fully efficient in all other contingencies. This inefficient equilibrium always exists under our assumptions on parameters. As a result, the worst-case welfare under this protocol is

$$\underline{\Lambda}_{\text{pre}} = 1 - \alpha(1-q)^2(1 - \underline{\lambda}_u(\text{pre})),$$

where α was defined in (B.1). To compare it to the least efficient equilibrium in the two post-trade transparent regimes, we can argue something stronger than stated in the Proposition which is that the pre-trade transparent market dominates the other two *conditional on Z*. This is clear for the case in which $Z = 1$, since then $\underline{\Lambda}_{\text{pre}}(1) = 1$. Conditional on $Z = 0$,

$$\underline{\Lambda}_{\text{pre}}(0) - \underline{\Lambda}_{\text{post}_{\text{high}}}(0) = \underline{\Lambda}_{\text{pre}}(0) - \underline{\Lambda}_{\text{post}_{\text{low}}}(0) = 1 - (1-q)^2(1 - \underline{\lambda}_u(\text{pre})) - \frac{q}{2} > 1 - (1-q)^2 - \frac{q}{2} > 0.$$

Next, we compare the regimes in terms to their most efficient equilibrium. If condition (B.5) is satisfied, then the pre-trade transparent market is fully efficient ($\bar{\Lambda}_{\text{pre}} = 1$) and the result follows immediately. Suppose that (B.5) is violated. Then, the most efficient equilibrium of the pre-trade transparent market is the one where the uninformed dealer quotes Δ in period zero and, conditional on a pair of quotes (Δ, Δ) , the customer agrees to trade with probability $\bar{\lambda}_u(\text{pre})$.

Consider first the low-search-cost regime. We will show that the best equilibrium in

the pre-trade transparent market results in higher welfare for every realization of Z . Given $Z = 1$, the pre-trade transparent market is fully efficient, and hence $\bar{\Lambda}_{\text{pre}}(1) = 1 > \bar{\Lambda}_{\text{post}_{\text{low}}}(1)$. To compare the regimes conditional on $Z = 0$, we write

$$\bar{\Lambda}_{\text{pre}}(0) = 1 - (1 - q)^2(1 - \bar{\lambda}_u(\text{pre})), \quad \bar{\Lambda}_{\text{post}_{\text{low}}}(0) \leq 1 - (1 - q)(1 - \bar{\lambda}_u(\text{post}_{\text{low}})),$$

where the expression for $\bar{\Lambda}_{\text{post}_{\text{low}}}(0)$ is written as a weak inequality because an equilibrium with search need not exist, in which case $\bar{\Lambda}_{\text{post}_{\text{low}}}(0) = \frac{q}{2} < 1 - (1 - q)(1 - \bar{\lambda}_u(\text{post}_{\text{low}}))$. So it suffices to show that $(1 - q)(1 - \bar{\lambda}_u(\text{pre})) \leq 1 - \bar{\lambda}_u(\text{post}_{\text{low}})$. To do this, write

$$\begin{aligned} 1 - \bar{\lambda}_u(\text{post}_{\text{low}}) - (1 - q)(1 - \bar{\lambda}_u(\text{pre})) &= \frac{1 - \delta\gamma\Delta}{1 + (1 - \delta\gamma/2)\Delta} - \frac{(1 - q)(1 - \Delta + \delta\gamma\kappa)}{1 + \Delta} \\ &\geq \frac{(1 - q)\Delta + q - \delta\gamma(\Delta + (1 - q)\kappa)}{1 + \Delta} > 0, \end{aligned}$$

where the strict inequality follows from assumption (3.1).

Now, let us carry out the comparison for the high-search-cost post-trade transparent market. As before, we have $\bar{\Lambda}_{\text{pre}}(1) = 1 \geq \bar{\Lambda}_{\text{post}_{\text{high}}}(1)$. Moreover, letting $\bar{\lambda}_u(\text{post}_{\text{high}})$ and $\bar{\lambda}_h(\text{post}_{\text{high}})$ be part of a solution to the program ($\text{Opt}_{\text{post}_{\text{high}}}$), we can write

$$\begin{aligned} \bar{\Lambda}_{\text{post}_{\text{high}}}(0) &= 1 - \frac{q}{2}(1 - \bar{\lambda}_h(\text{post}_{\text{high}})) - (1 - q)(1 - \bar{\lambda}_u(\text{post}_{\text{high}})) \\ &\leq 1 - (1 - q)(1 - \bar{\lambda}_u(\text{post}_{\text{high}})) \leq 1 - (1 - q)(1 - \bar{\lambda}_u(\text{post}_{\text{low}})) < \bar{\Lambda}_{\text{pre}}(0), \end{aligned}$$

where the second inequality follows from (IC1(0)), which implies that $\bar{\lambda}_u(\text{post}_{\text{high}}) \leq \bar{\lambda}_u(\text{post}_{\text{low}})$.

Finally, we establish the second part of Proposition 9. Suppose that $q \geq \frac{\bar{\lambda}_u(\text{post}_{\text{low}}) - \underline{\lambda}_u(\text{pre})}{1 - \underline{\lambda}_u(\text{pre})}$. When $Z = 1$, we have $\underline{\Lambda}_{\text{pre}}(1) = 1 \geq \max\{\bar{\Lambda}_{\text{post}_{\text{low}}}(1), \bar{\Lambda}_{\text{post}_{\text{high}}}(1)\}$. When $Z = 0$, we showed above that

$$\max\{\bar{\Lambda}_{\text{post}_{\text{low}}}(0), \bar{\Lambda}_{\text{post}_{\text{high}}}(0)\} \leq 1 - (1 - q)(1 - \bar{\lambda}_u(\text{post}_{\text{low}})) \leq 1 - (1 - q)^2(1 - \underline{\lambda}_u(\text{pre})),$$

where the second inequality follows from the assumption on q . Hence, $\bar{\Lambda}_{\text{pre}} > \underline{\Lambda}_{\text{pre}} \geq \max\{\bar{\Lambda}_{\text{post}_{\text{low}}}, \bar{\Lambda}_{\text{post}_{\text{high}}}\}$ when q is high enough.

Proof of Proposition 10

Step 1: Continuation equilibrium in period one. As in the pre-trade transparent market, it is still the case that dealers learn each other's types by observing the quotes provided in period zero. The only difference, relative to that regime, is that now the period-one customer will enter the market with a belief about v different from the prior (since it will be refined by the

outcome of trade in period zero). Still, the arguments that we gave in the proof of Proposition 8 for showing that the Bertrand outcome is the unique continuation equilibrium in period one apply here as well since they only relied on the lack of information asymmetry between dealers and were, in particular, independent of the beliefs of the period-one customer at the time of entering the market.

Step 2: Equilibrium strategies in period zero.

2.a) Necessary conditions for equilibrium strategies. Again, applying the arguments from Proposition 8, it must be that the structure of dealers' strategies coincides with the one in the pre-trade transparent market. Namely, the low and high type dealers will follow exactly the same quoting strategy as in the pre-trade transparent regime, and the uninformed dealer will play a pure strategy on $b_u \leq \Delta$. However, as we show next, the introduction of post-trade transparency will affect dealers' continuation payoffs conditional on deviating in period zero, and hence will have an effect on trading outcomes through dealers' incentive-compatibility constraints. More precisely, the equilibrium values of b_u and the probability that the customer agrees to trade when observing a pair of quotes equal to (Δ, Δ) will be different relative to the regime without post-trade transparency.

2.b) Construction of a fully efficient equilibrium. The condition preventing the uninformed type from deviating to 1 is unchanged, and hence we must have $b_u \geq \underline{b}_u(\text{pre})$. Consider the low-type dealer. Her on-path payoffs are the same as those in the pre-trade transparent market, so we only need to compute her payoff from deviating to b_u in period zero. If the opponent's type is l , the winning quote will reveal that $v = -1$ and the continuation payoff is 0. If the opponent's type is u , there will be trade at a price equal to b_u and hence the period-one customer will enter the market with beliefs that assign probability zero to type l . Because of the never-dissuaded-once-convinced refinement, this in turn implies that, after deviating in period zero, the low-type dealer can undercut her opponent's quote (which is equal to 0) in the following period without facing belief-based punishment, which gives her a payoff arbitrarily close to 1. As a result, her incentive constraint in period zero becomes

$$(1 - q)\Delta \geq \frac{1 - q}{2}(b_u + 1) + \delta\gamma(1 - q) \iff b_u \leq 2(\Delta - \delta\gamma) - 1 \equiv \bar{b}_u(\text{full}).$$

Observe that $\bar{b}_u(\text{full}) < \bar{b}_u(\text{pre})$, which in turn implies that we have a more stringent condition for existence of an efficient equilibrium in the fully transparent market, given by

$$\underline{b}_u(\text{pre}) \leq \bar{b}_u(\text{full}) \iff \delta\gamma \leq \frac{(1 - q)(2\Delta - 1) - q}{3(1 - q)} \equiv \bar{\delta}_{\text{full}} < \bar{\delta}_{\text{pre}}.$$

2.c) Construction of an inefficient equilibrium. Next, consider equilibria in which, in period

zero, the uninformed dealer quotes $b_u = \Delta$ and the customer agrees to trade with probability λ_u after observing a pair of quotes equal to (Δ, Δ) . As in the previous case, the incentive-compatibility constraint of the uninformed dealer is the same as in the pre-trade transparent market, which yielded the lower bound $\lambda_u \geq \underline{\lambda}_u(\text{pre})$. For the low-type dealer, the continuation payoff if she deviates to Δ in period zero is again equal to 1. Thus, we have

$$(1 - q)\Delta \geq \frac{(1 - q)\lambda_u}{2}(\Delta + 1) + \delta\gamma(1 - q) \iff \lambda_u \leq \frac{2(\Delta - \delta\gamma)}{1 + \Delta} \equiv \bar{\lambda}_u(\text{full}).$$

Observe that $\bar{\lambda}_u(\text{full}) < \bar{\lambda}_u(\text{pre})$ as stated in the Proposition.

Finally, an equilibrium in which the uninformed dealer quotes Δ in period zero exists if and only if

$$\underline{\lambda}_u(\text{pre}) \leq \bar{\lambda}_u(\text{full}) \iff \delta\gamma \leq \frac{(2 - q)\Delta^2 - q}{(1 - q)(1 + 3\Delta)} = \bar{\delta}_{\text{exist}},$$

which is precisely our assumption (3.1).

The above conditions are also sufficient for an equilibrium if we set off-path beliefs in period zero in the same way as in the proof of Proposition 8.

Proof of Proposition 11

First, both regimes share the same least efficient equilibrium, with welfare

$$\underline{\Lambda}_{\text{pre}} = \underline{\Lambda}_{\text{full}} = \alpha(1 - (1 - q)^2(1 - \underline{\lambda}_u(\text{pre}))) + 1 - \alpha,$$

where α was defined in (B.1). Therefore, it suffices to compare the two regimes in terms of their most efficient equilibrium.

If $\delta\gamma \leq \bar{\delta}_{\text{full}} < \bar{\delta}_{\text{pre}}$, then both regimes have a fully efficient equilibrium. If $\delta\gamma \in (\bar{\delta}_{\text{full}}, \bar{\delta}_{\text{pre}}]$, then only the pre-trade transparent market has a fully efficient equilibrium, and hence $\bar{\Lambda}_{\text{pre}} = 1 > \bar{\Lambda}_{\text{full}}$. Otherwise, if $\delta\gamma > \bar{\delta}_{\text{pre}}$, then neither regime has a fully efficient equilibrium. The respective upper bounds on equilibrium welfare are given by

$$\bar{\Lambda}_{\text{pre}} = \alpha(1 - (1 - q)^2(1 - \bar{\lambda}_u(\text{pre}))) + 1 - \alpha, \quad \bar{\Lambda}_{\text{full}} = \alpha(1 - (1 - q)^2(1 - \bar{\lambda}_u(\text{full}))) + 1 - \alpha.$$

Since $\bar{\lambda}_u(\text{pre}) > \bar{\lambda}_u(\text{full})$, it follows that $\bar{\Lambda}_{\text{pre}} > \bar{\Lambda}_{\text{full}}$. Hence, the pre-trade transparent regime weakly welfare-dominates the fully transparent one if and only if $\delta\gamma > \bar{\delta}_{\text{full}}$.

Proof of Proposition 12

Using the discounted probabilities derived in Lemma 1, the optimal mechanism maximizes

$$\mathbb{E} \left[\sum_{i \in \{A, B\}} (x_0(\theta_i, \theta_{-i}) + \delta\gamma x_1(\theta_i, \theta_{-i})) \right] \quad (Opt_{mech})$$

over all feasible mechanisms, as defined in Section 4.

Step 1: Optimality of static mechanisms. Let (x^*, p^*, b^*) be an optimal *static* mechanism, that is, optimal for (Opt_{mech}) subject to the additional restriction of being constant in Z . We first show that (x^*, p^*, b^*) is also a solution to (Opt_{mech}) . Suppose otherwise. Let $(x_Z, p_Z, b_Z)_{Z \in \{0, 1\}}$ be a feasible (non-static) mechanism satisfying

$$\mathbb{E} \left[\sum_{i \in \{A, B\}} (x_0(\theta_i, \theta_{-i}) + \delta\gamma x_1(\theta_i, \theta_{-i})) \right] > (1 + \delta\gamma) \mathbb{E} \left[\sum_{i \in \{A, B\}} x^*(\theta_i, \theta_{-i}) \right].$$

Now, consider the following static mechanism: $(\hat{x}, \hat{p}, \hat{b}) = \left(\frac{x_0 + \delta\gamma x_1}{1 + \delta\gamma}, \frac{p_0 + \delta\gamma p_1}{1 + \delta\gamma}, \frac{b_0 + \delta\gamma b_1}{1 + \delta\gamma} \right)$. The fact that $(x_Z, p_Z, b_Z)_{Z \in \{0, 1\}}$ is feasible in (Opt_{mech}) implies that so is $(\hat{x}, \hat{p}, \hat{b})$. Moreover, the value of the objective at $(\hat{x}, \hat{p}, \hat{b})$ is

$$\mathbb{E} \left[\sum_{i \in \{A, B\}} x_0(\theta_i, \theta_{-i}) + \delta\gamma x_1(\theta_i, \theta_{-i}) \right] > (1 + \delta\gamma) \mathbb{E} \left[\sum_{i \in \{A, B\}} x^*(\theta_i, \theta_{-i}) \right],$$

contradicting optimality of (x^*, p^*, b^*) among static mechanisms.

Step 2: Solving for the optimal static mechanism. Given that we can focus on static mechanisms, to simplify notation, we drop the Z subscripts from the problem. Suppose first that $\Delta \geq 1/2$. In that case, the mechanism proposed in the proposition is fully efficient and hence cannot be improved upon. One can also readily check that it is feasible.

Next, suppose that $\Delta < 1/2$. We can then show that any implementable mechanism satisfies $x(0, 0) \leq \Delta$. To do so, combine the ex-post IC constraint of the low type when her opponent is uninformed with the ex-post IR constraint of the uninformed type when her opponent is uninformed to obtain that $x(0, 0) \leq x(-1, 0) + p(-1, 0)$. On the other hand,

$$\Delta \geq x(-1, 0) + x(0, -1) + b(0, -1) \geq x(-1, 0) + x(0, -1) + p(-1, 0) + p(0, -1) \geq x(-1, 0) + p(-1, 0),$$

where the first inequality comes from buyer's ex-post IR constraint, the second one from the designer's ex-post budget-balance constraint, and the third one from the dealer's ex-post IR constraint, all evaluated at the type profile $(0, -1)$. These two conditions combined imply

$x(0, 0) \leq \Delta$. But this implies that welfare under any implementable mechanism is bounded above by $\Lambda^* \equiv 1 - (1 - q)^2(1 - 2\Delta)$, i.e., welfare attained when trade happens with probability one under all contingencies, except when both dealers are uninformed and each trades with probability Δ . Our proposed mechanism is feasible and attains precisely that bound.

Step 3: Comparison with the pre-trade transparent market. The comparison with respect to the most efficient equilibrium in the pre-trade transparent market is trivial when $\Delta \geq 1/2$. If $\Delta < 1/2$, note that $\bar{\delta}_{\text{pre}} < 0$, and thus the pre-trade transparent regime fails to have a fully efficient equilibrium. Thus, welfare in the most efficient equilibrium is given by $\bar{\Lambda}_{\text{pre}} = 1 - \alpha(1 - q)^2(1 - \bar{\lambda}_u(\text{pre}))$. The welfare difference is

$$\Lambda^* - \bar{\Lambda}_{\text{pre}} = \frac{(1 - q)^2(2\Delta^2 + \delta\gamma\Delta(1 + 2\Delta))}{(1 + \Delta)(1 + \delta\gamma)} > 0.$$

Proof of Proposition 13

Step 1: Specifying equilibrium strategies. Because there are only two dealers, their strategies in the auction can be described by the price at which they quit the auction, conditional on the opponent not having done so already. We refer to this price as a dealer's "quote". We construct an ex-post equilibrium of the game where, in period zero, dealer i with type θ quotes θ . In period one, dealer i with type 0 quotes 0 if she learned that her opponent has a high type, and quotes $-1 + \Delta$ if she learned that her opponent has a low type in the previous period. In all other cases, dealer i with type θ quotes θ in period one. Conditional on the opponent having quoted 1 in period zero and not having quit at price 1 in period one, we set the uninformed dealer's off-path belief to assign probability one to the opponent being uninformed.

The customer agrees to trade with probability one at the end of the auction when she is indifferent in all cases except after she observes that both dealers quoted 0, in which case she accepts with probability $2 \min\{\Delta, 1/2\}$ (with the tie broken uniformly across the two dealers) and exits otherwise. The only on-path case in which she strictly prefers to exit or request another quote over accepting is when only one dealer quotes 0, or when only one dealer quotes $-1 + \Delta$, which only happens when one dealer is uninformed and the other one has type -1 . In such a case, we set her strategy to be to ask for another quote with probability one. Conditional on the customer having requested another quote, the winning dealer's strategy is to counteroffer $\hat{v} + \Delta$, where \hat{v} is the customer's estimate of the value at this stage, and the customer accepts with probability one. These strategies implement precisely the allocation from Proposition 12.

Step 2: Checking players' incentive constraints. Let us now check that this is an equilibrium

by solving the game backwards. The strategies of the customer and the winning dealer are clearly optimal in the node of the game in which the customer requests another quote. Let us check dealers' incentive constraints in the auction. The high-type dealer gets a zero payoff no matter what her strategy is and regardless of the opponent's type. The same is true for the uninformed dealer in period one after having played her equilibrium strategy in period zero. In period zero, by bidding 0 and playing her continuation strategy, the uninformed dealer gets a zero payoff. If she deviates by quoting a lower price she also gets 0, regardless of the opponent's type. The only relevant deviation is then to quote 1. If the opponent's type is either low or high, this deviation leads to zero profits in both periods. If the opponent is uninformed, dealer i loses the auction in period zero and gets a zero payoff. Moreover, if i does not quit at price 1 in period one, the opponent will stay until price 0, and hence the continuation payoff is 0 as well. Thus, this deviation is not profitable either. Finally, we look at the low type. If the opponent is uninformed, dealer i 's payoff is Δ in periods zero and one. This is because, in period one, the uninformed opponent will quote $-1 + \Delta$ and, in period zero, the opponent will quote 0. In both cases, the customer will request another quote. Dealer i will then counteroffer $-1 + \Delta$, which the customer will accept with probability one. If dealer i instead quotes 0, then the opponent will believe that she is uninformed and will quote 0 in period one. Dealer i can then ensure a continuation payoff of $\min\{\Delta, 1/2\}$ if she quotes 0 in period one, or at most Δ if she does not quit at price 0 and wins the auction with probability one. In period zero, she would get $\min\{\Delta, 1/2\}$. Hence, the payoff from quoting -1 dominates the payoff from quoting 0. Any quote higher than 0 is also suboptimal. Finally, if the opponent has a low type, any quote leads to a zero payoff, and hence the incentive constraint is trivially satisfied.

Online Appendix

OA.1 Two-sided private information

In this section, we consider a perturbed version of our model in which dealers face a small degree of uncertainty about customers' private value for the asset. Formally, suppose that the period- t customer's payoff from trading the asset is $v_t + \Delta + \epsilon_t$, where ϵ_t is uniformly distributed on $[-\underline{\epsilon}, \underline{\epsilon}]$, independently across t and of v_t . The parameter $\underline{\epsilon} > 0$ is a commonly known constant, which should be thought of as small. We assume that ϵ_t is privately observed by the period- t customer at the time of entering the market.

It turns out that introducing a small amount of customer private information only influences equilibrium trading outcomes through customers' directed search behavior in the post-trade transparent market. In particular, in all of the regimes, dealers' quoting strategies as well as the probability that on-path quotes are accepted are virtually unaffected by this perturbation. The only adjustment that is required in the proofs is that the interior probabilities of trade that are required to satisfy dealers' incentive compatibility are now implemented by slightly modifying dealers' quotes. This modification in the quotes is constructed so as to ensure that there is a cutoff type in $[-\underline{\epsilon}, \underline{\epsilon}]$ such that all customers above the cutoff agree to trade with probability one. In this way, the extension allows to purify the random acceptance decision by the customer that is a key aspect driving inefficiency in all the trading regimes that we study.

In order to highlight the role that uncertainty about the customer's private value has on directed search, we introduce and study an additional equilibrium refinement that we call *search symmetry*: In case of indifference about which dealer to contact first, the customer must randomize with uniform probabilities. What makes search symmetry plausible is the realistic possibility that customers have a "default dealer" whom they always visit first, except when they can strictly benefit from doing otherwise (assuming each dealer is ex-ante equally likely to be the "default dealer").

Proposition OA.1. *There exists $\underline{\epsilon}^* > 0$ such that, when $\underline{\epsilon} < \underline{\epsilon}^*$,*

- 1. Under search symmetry, there exists a non-empty open interval of values of q for which the low-search-cost post-trade transparent market cannot be welfare-ranked with respect to its opaque counterpart: It has a strictly better worst-case equilibrium but a strictly worse best-case equilibrium.*
- 2. With or without the search symmetry refinement, there exists a non-empty open interval of values of q for which the high-search-cost post-trade transparent market cannot be*

welfare-ranked with respect to its opaque counterpart: It has a strictly better worst-case equilibrium but a strictly worse best-case equilibrium.

3. All of the remaining welfare rankings coincide with those under $\underline{\epsilon} = 0$ (regardless of whether search symmetry is assumed or not).

Proof. Because welfare in all other regimes remains unchanged,⁴² in order to avoid repetition, we focus on discussing which trading outcomes are affected by this perturbation in the post-trade transparent market. Throughout, we rely on the results and notation from Appendix B. Part 1. Throughout this part of the proof, we take s to be small enough.⁴³ The least efficient equilibrium remains unchanged. Moreover, provided that a search equilibrium exists, the most efficient equilibrium must feature search in period zero. In period one, following trade at a fully informative price, the continuation outcome remains the no-search one where the dealer who gets contacted behaves as a monopolist. To derive dealers' optimal monopoly price, observe that when dealers' (commonly known) value is $v \in \{-1, 1\}$, their profits from quoting a price $p \in [v + \Delta - \underline{\epsilon}, v + \Delta + \underline{\epsilon}]$ is

$$(p - v)(1 - F(p - v - \Delta)),$$

where F is the cdf of ϵ_t . Profits are strictly decreasing in $p \in (v + \Delta - \underline{\epsilon}, v + \Delta + \underline{\epsilon})$ when $\underline{\epsilon}$ is sufficiently small. Therefore, the optimal monopoly price is $p = v + \Delta - \underline{\epsilon}$, and the resulting trading outcome is efficient. Because both dealers quote the same price in equilibrium, the customer in period one is indifferent and randomizes by visiting each of them with probability 1/2 under the search symmetry refinement. The outcome following no trade in period zero is unaffected as well.

Following trade at the uninformed type's price, it must be the case that the dealer who traded in period zero, say dealer A , provides a quote of at least $1 + \Delta$ and gets visited with zero probability. Dealer B , on the other hand, trades according to the static monopoly trading pattern. In particular, the low type trades with probability one at a price equal to $-1 + \Delta - \underline{\epsilon}$, and the uninformed and high types provide a quote equal to b_u and b_h , respectively, so as to ensure the following incentive-compatibility conditions

$$\Delta \geq (b_u + 1)(1 - F(b_u - \Delta)), \quad b_u(1 - F(b_u - \Delta)) \geq b_h(1 - F(b_h - 1 - \Delta)).$$

Observe that b_u and b_h have to be arbitrarily close to, respectively, Δ and $1 + \Delta$, when $\underline{\epsilon}$ is small if the uninformed and high types trade with positive probability. Thus, the respective

⁴²Note that prices do change for any strictly positive $\underline{\epsilon}$; however, the probabilities of trade are not affected, since these probabilities are pinned down by the dealers' incentive constraints.

⁴³It will become clear that the upper bound on s can be taken to be independent of the bound on $\underline{\epsilon}$.

upper bounds on their trade probabilities are arbitrarily close to those obtained under $\underline{\epsilon} = 0$. These two types cannot profitably deviate by lowering their quotes if we set off-path beliefs so that lower quotes get rejected. Moreover, because their profit is strictly decreasing in price, they do not benefit from deviating to a higher quote either. Given dealers' quoting strategies at this node of the game, the period-one customer strictly prefers to contact dealer B first regardless of her private type ϵ . This is because she makes a strictly positive payoff at least in the event that dealer B has a low type. Moreover, she will never find it worthwhile to pay the search cost (no matter how small) and contact dealer A second, thus ensuring that dealer A never gets contacted as required.

Having derived continuation play in period one, we can write incentive-compatibility constraints in period zero as in the proof of Proposition 5. The only difference is that, whenever imposed, the restriction on the customer's directed search strategy in period one now pins down the probabilities that dealers get contacted after trading at a fully revealing quote in period zero. As a result, the low-type's incentive constraint becomes

$$\Delta + \delta\gamma\frac{1}{2}\Delta \geq \lambda_u(1 + \Delta) + \delta\frac{\gamma(1 - \lambda_u)}{2}\Delta,$$

where λ_u is the probability that the uninformed type trades in period zero (which pins down her equilibrium quote). Similarly, for the uninformed type we have

$$\lambda_u\Delta + \frac{q}{2}(\Delta - 1) \geq \frac{q}{2}\Delta + \frac{\delta\gamma}{2}\left(\frac{q}{2}\Delta + (1 - q)b_u(1 - F(b_u - \Delta))\right).$$

Combining everything, the condition for the existence of a search equilibrium when $\underline{\epsilon}$ is small can be written as

$$q \leq \frac{2\Delta^2}{(1 + (1 - \frac{\delta\gamma}{2})\Delta)(1 + \frac{\delta\gamma}{2}\Delta)}.$$

This condition is more stringent than the condition derived for the low-search-cost opaque market. The reason is that, once we shut down the possibility of using the customer's directed search strategy in period one to provide incentives, the low-type's incentive constraint is unambiguously more stringent in the post-trade transparent market (relative to the opaque one) due to the signaling motive. As a result, different from Proposition 6, there is now a region of intermediate values of q under which the post-trade transparent regime will have a strictly worse best-case equilibrium compared to the opaque one. Thus, the two regimes cannot be welfare-ranked in that region.

Part 2. Let us again begin by deriving continuation outcomes in period one. Under high

search costs, it is also the case that the only instance in which the period-one customer may not be indifferent regarding which dealer to contact first is after trade occurs at the uninformed dealer's quote in period zero. As argued in the proof of Proposition 4, in that case, dealer A (the dealer who traded in period zero) must trade with probability one conditional on being contacted, and will therefore provide a quote equal to $\Delta - \underline{\epsilon}$. Dealer B , on the other hand, must trade with interior probability in the event of being uninformed or high-type. Because this interior probability of trade is sustained by having dealer B quote a higher price, it follows that the customer will strictly prefer to contact dealer A first regardless of her private type.

The above implies that the incentive constraint of the low type in period zero is more stringent in the perturbed version of the game, and it will be given by

$$\Delta + \frac{\delta\gamma}{2}\Delta \geq \lambda_u(1 + \Delta) + \frac{\delta\gamma}{2}(2\lambda_u(1 + \Delta) + (1 - \lambda_u)\Delta). \quad (\text{OA.1})$$

As in the low-search-cost environment, we can argue that, under the perturbation, it is no longer possible to rank the post-trade transparent and the opaque market under certain values of q . In particular, taking q to be sufficiently low, condition (OA.1) will bind in the most efficient equilibrium. In fact, fixing a small value of $\underline{\epsilon}$ and taking the limit as $q \rightarrow 0$ we obtain that the best-case welfare difference between the opaque and the post-trade transparent market in the $\underline{\epsilon}$ -perturbed game satisfies

$$\lim_{q \rightarrow 0} (\bar{\Lambda}_{\text{opaq high}}^{\underline{\epsilon}} - \bar{\Lambda}_{\text{post high}}^{\underline{\epsilon}}) \approx \frac{\delta\gamma\Delta^2}{(1 + \Delta)(2(1 + \Delta) + \delta\gamma(2 + \Delta))} > 0.$$

Thus, there is again an open set of values of q under which the post-trade transparent and the opaque regimes cannot be welfare-ranked.

Part 3. One can verify that Proposition 7 still holds when $\underline{\epsilon} > 0$ is small enough. Finally, since all the changes that arise in this section go in the direction of making the best-case equilibrium in the post-trade transparent market weakly worse, none of the other welfare rankings in the paper are affected by it. So, overall, this perturbation of the model only affects the results in that it invalidates the result in Proposition 6 under some values of q . \square

OA.2 Equilibrium refinements

D1 equilibrium in the static opaque market

In this section, we apply the divinity criterion (D1) by Banks and Sobel (1987) to the opaque market with high search cost. Intuitively, the divinity criterion states that an off-path belief

should only place probability weight on types of opponents for which the deviation could lead to a strict improvement over their equilibrium payoff. We show that this refinement uniquely selects the most efficient separating equilibrium.

In the high-search-cost setting, the customer does not search in equilibrium, so we can restrict attention to the two-player game that ensues between the customer and the dealer that she randomly matches with. The customer's strategy can be described by $a \in \{0, 1\}$, which specifies whether or not she agrees to trade. Given a quote $b \in \mathbb{R}$, an acceptance decision by the customer, and the dealer's type $\theta \in \{-1, 0, 1\}$, the dealer's and the customer's payoffs can be written as $u_D(a, b, \theta) = a(b - \theta)$ and $u_C(a, b, \theta) = a(\theta + \Delta - b)$, respectively. Let $\pi : \mathbb{R} \rightarrow \Delta\{-1, 0, 1\}$ be the customer's posterior belief about θ after observing the dealer's quote, and let $\text{MBR}(\pi, b) \subset [0, 1]$ be the set of acceptance probabilities that are a best reply for the customer, given the quote b and her belief π . Let

$$D(\theta, b) \equiv \bigcup_{\pi \in \Delta\{-1, 0, 1\}} \{\beta \in \text{MBR}(\pi, b) : u_D^*(\theta) < u_D(\beta, b, \theta)\},$$

where $u_D^*(\theta)$ is the equilibrium payoff of dealer type θ . Let $\bar{D}(\theta, b)$ be the closure of $D(\theta, b)$. The D1 criterion requires that the customer's beliefs, following an off-path quote b , assign zero probability to type θ whenever there exists θ' such that $\bar{D}(\theta, b) \subsetneq D(\theta', b)$.

Proposition OA.2. *In the high-search-cost opaque market, the unique D1 Perfect Bayesian equilibrium coincides with the best equilibrium under the solution concept defined in Section 2.*

Proof. Step 1: Ruling out pooling equilibrium. Suppose by contradiction that there are types $\theta < \theta'$ and a price b^* that belongs to the support of the equilibrium strategy of both types θ and θ' . Equilibrium payoffs for these two types are $u_D^*(\theta) = a(b^*)(b^* - \theta)$ and $u_D^*(\theta') = a(b^*)(b^* - \theta')$, where $a(b^*)$ is the equilibrium acceptance probability of the customer following b^* . Consider first the case in which $a(b^*) > 0$. By customer's optimality, this requires that $b^* < \theta' + \Delta$. Now, consider a quote $\tilde{b} = \theta' + \Delta - \epsilon$, where $\epsilon > 0$ is arbitrarily small. We argue that such a quote must be off-the-equilibrium-path. First, \tilde{b} cannot be offered only by types weakly larger than θ' because, if it were, it would be accepted by the customer with probability one and type θ would benefit from deviating to it. Therefore, some type $\hat{\theta} < \theta'$ must be making this offer in equilibrium. But since no type $\theta'' > \theta'$ would be willing to offer $\theta' + \Delta - \epsilon < \theta''$, such a quote gets rejected with probability one by the customer. This is a contradiction, since type $\hat{\theta}$ can ensure a strictly positive payoff by offering b^* instead.

Since \tilde{b} is an off-path quote, we apply the D1 criterion to it, and show that beliefs following \tilde{b} assign probability one to types θ' and higher. To see this, applying the above definitions,

we have

$$\bar{D}(\theta, \tilde{b}) = \left[\frac{a(b^*)(b^* - \theta)}{\tilde{b} - \theta}, 1 \right], \quad D(\theta', \tilde{b}) = \left(\frac{a(b^*)(b^* - \theta')}{\tilde{b} - \theta'}, 1 \right).$$

When ϵ is sufficiently small, the fact that $a(b^*) > 0$ and $\theta' > \theta$ implies that

$$\frac{a(b^*)(b^* - \theta)}{\tilde{b} - \theta} > \frac{a(b^*)(b^* - \theta')}{\tilde{b} - \theta'},$$

and therefore $\bar{D}(\theta, \tilde{b}) \subsetneq D(\theta', \tilde{b})$. This implies that type θ does not belong to the support of $\pi(\tilde{b})$. Now, consider any other $\hat{\theta} < \theta$ (in our setting we three types, this implies $\hat{\theta} = -1$, $\theta = 0$ and $\theta' = 1$). By incentive compatibility, it must be that $u_D^*(\hat{\theta}) \geq a(b^*)(b^* - \hat{\theta})$, and thus $\bar{D}(\hat{\theta}, \tilde{b}) \subset \left[\frac{a(b^*)(b^* - \hat{\theta})}{\tilde{b} - \hat{\theta}}, 1 \right]$. An analogous argument then shows that $\bar{D}(\hat{\theta}, \tilde{b}) \subsetneq D(\theta', \tilde{b})$, thus excluding type $\hat{\theta}$ from the support of $\pi(\tilde{b})$ as well. But then, if the customer's off-path beliefs following $\theta' + \Delta - \epsilon$ are fully supported on types θ' and higher, she will strictly want to accept such a quote. This is a contradiction, since then types θ and θ' would benefit from deviating to \tilde{b} .

It remains to consider the case with $a(b^*) = 0$. In this case, types θ and θ' get zero equilibrium payoffs. This implies that $\theta = 0$ and $\theta' = 1$, since type -1 can always ensure a payoff of Δ by quoting a price close to $-1 + \Delta$. Consider a quote $\tilde{b} = \Delta - \epsilon$, when $\epsilon > 0$ is small enough. By the same arguments as in the previous case, such a quote must be off-path. We now argue that D1 requires the belief $\pi(\tilde{b})$ to assign zero probability to type -1 . To see this, note that $u_D^*(-1) \geq \Delta$, and therefore $\bar{D}(-1, \tilde{b}) \subset \left[\frac{\Delta}{1 + \tilde{b}}, 1 \right]$. On the other hand, since $a(b^*) = 0$, we have $D(0, \tilde{b}) = (0, 1]$. Thus, $\bar{D}(-1, \tilde{b}) \subsetneq D(0, \tilde{b})$. Hence, the customer's belief following \tilde{b} assigns zero probability to type -1 . This yields a contradiction, since then the customer would accept this quote with probability one, giving rise to a profitable deviation for type 0 whose equilibrium payoff is 0.

Step 2: Applying D1 to the set of separating equilibria. In any separating equilibrium, $u_D^*(-1) = \Delta$, $u_D^*(0) = \lambda_0 \Delta$ and $u_D^*(1) = \lambda_1 \Delta$, where λ_0 and λ_1 are the respective probabilities that the customer agrees to trade at a price equal to Δ and $1 + \Delta$, respectively. Consider a quote $\tilde{b} = \Delta - \epsilon$ where $\epsilon > 0$ is arbitrarily small. Applying the above definitions, we have

$$\bar{D}(-1, \tilde{b}) = \left[\frac{\Delta}{\tilde{b} + 1}, 1 \right], \quad D(0, \tilde{b}) = \left(\frac{\lambda_0 \Delta}{\tilde{b}}, 1 \right).$$

In order to sustain a fully separating equilibrium, it cannot be that the customer's beliefs

after observing $\tilde{b} = \Delta - \epsilon$, for ϵ small enough, rule out the low type. Otherwise, the low and the uninformed types could trade with probability one by quoting \tilde{b} . For this to be consistent with D1, $\bar{D}(-1, \tilde{b}) \not\subseteq D(0, \tilde{b})$ could not hold, which is the case if and only if

$$\frac{\Delta}{\tilde{b} + 1} \leq \frac{\lambda_0 \Delta}{\tilde{b}}.$$

Taking ϵ to 0 yields $\lambda_0 \geq \Delta/(1 + \Delta)$.

An analogous argument shows that $\lambda_1 \geq \lambda_0 \Delta/(1 + \Delta)$ in any D1 equilibrium. Since we have already shown in the proof of Proposition 1 that any separating equilibrium must satisfy $\lambda_0 \leq \Delta/(1 + \Delta)$ and $\lambda_1 \leq \lambda_0 \Delta/(1 + \Delta)$, it follows that the D1 refinement uniquely selects the most efficient separating equilibrium. One can also check that these equilibrium strategies, together with the off-path beliefs that we consider throughout the paper, indeed form a D1 equilibrium. \square

Refinements on customer's beliefs

In this section, we discuss some implausible equilibria that would arise if we were to dispense of either of the two refinements that we impose on the customer's off-path beliefs. We start by considering the never-dissuaded-once-convinced axiom. As we argued in the paper, this axiom ensures that, in the regimes with post-trade transparency, the continuation outcome after a trade that fully reveals the state is the efficient one. We show that, without the axiom, it is possible to construct an equilibrium in the post-trade transparent market where the continuation outcome after a fully revealing quote is inefficient. For concreteness, we do this for the high-search-cost setting, although a similar construction obtains under low search costs as well.

Claim OA.1. *Without the never-dissuaded-once-convinced axiom, the high-search-cost post-trade transparent market admits an equilibrium in which trade happens with zero probability following a quote that (publicly) reveals the value of the asset to be high.*

Proof. Consider the following equilibrium construction. Dealers' quoting strategies at $Z_t = 0$ are the same as in the proof of Proposition 4. At $Z_t = 1$ and following a fully revealing trade at a price equal to $1 + \Delta$, we set dealers' strategies to be such that both the uninformed and the high type provide quotes strictly above $1 + \Delta$ and the customer always rejects. To ensure that this is a continuation equilibrium, we set the customer's off-path beliefs at this node of the game to assign probability one to type -1 after any off-path quote weakly below $1 + \Delta$. These off-path beliefs are obviously not consistent with our refinement, since they entail the customer revising her belief after learning that the state is high. In all other on-path $Z_t = 1$

histories, we set the continuation equilibrium to be the same as in Proposition 4. It is now straightforward to check that the chosen $Z_t = 0$ strategies and continuation equilibria form an equilibrium of the overall game. \square

Second, we discuss the role of the refinement on customer's beliefs in the pre-trade transparent regimes.

Claim OA.2. *Suppose that customers' beliefs are no longer required to be consistent with the on-path quote after observing it simultaneously with an off-path quote. We can then construct an equilibrium of the pre-trade transparent market where the uninformed and the high type (at all histories except at $Z_t = 1$ after the uninformed type learns that the state is -1) provide quotes strictly above $1 + \Delta$, which the customer always rejects.*

Proof. Fix the high type's and the uninformed type's quotes in both Z_t states to be $b_u > 1 + \Delta$ and $b_h > 1 + \Delta$, respectively, with $b_u < b_h$. Let $v(b_A, b_B)$ be the customer's posterior estimate of the asset's value following a pair of quotes (b_A, b_B) . We set these beliefs to satisfy

$$v(b_h, \tilde{b}) = \begin{cases} 1, & \text{if } \tilde{b} \geq b_u, \\ -1, & \text{if } \tilde{b} < b_u, \end{cases} \quad \text{and} \quad v(b_u, \tilde{b}) = \begin{cases} 1, & \text{if } \tilde{b} \geq b_h, \\ 0, & \text{if } \tilde{b} \in [b_u, b_h), \\ -1, & \text{if } \tilde{b} < b_u. \end{cases}$$

These beliefs are monotone and consistent with Bayes' rule on-path. They do not satisfy our refinement on beliefs, since it would require, for example, that $v(b_h, \tilde{b}) = 1$ for all off-path \tilde{b} .

Let us verify that this can be sustained as an equilibrium. If either one of the two dealers has a low type, the continuation equilibrium at $Z_t = 1$ is the Bertrand outcome described in Proposition 8. Otherwise, at $Z_t = 1$, the uninformed and the high type provide quotes $b_u, b_h > 1 + \Delta$ and never trade. They cannot profitably deviate, since the customer's off-path beliefs dictate that they can only trade by quoting a price less than $-1 + \Delta$ which yields a loss for these two types of dealers. The same argument implies that the uninformed and high types cannot profitably deviate at $Z_t = 0$. For the low type, we set her $Z_t = 0$ strategy to be the same as in Proposition 8. If she quotes a price higher than $-1 + \Delta$, she trades with zero probability at $Z_t = 0$, and her optimal continuation payoff is $(1 - q)\Delta$. Therefore, the condition preventing a deviation to a quote outside of the support of her equilibrium strategy is $(1 - q)\Delta \geq \delta\gamma(1 - q)\Delta$, which is always satisfied. \square

The equilibrium from Claim OA.2 is highly inefficient. In particular, it features zero trade in both Z_t states conditional on the asset value being high.

OA.3 Unique implementation of the optimal mechanism

In this appendix, we propose an indirect mechanism that uniquely implements the optimal direct mechanism from Proposition 12.

Consider the following trading mechanism, which we refer to as a *second-price auction with modified payments and rationing*. In every period t :

- Dealers submit a quote in $B \equiv \{-1, \underline{b}, \bar{b}, 0, 1\}$, where $\underline{b} \in (-1, -1 + \min\{\Delta, 1/2\})$, $\bar{b} \in (-1 + \Delta, 0)$.
- The period- t buyer observes dealers' quotes and decides whether or not to trade.
- Let b_i denote the i th lowest quote. If the buyer accepts, the platform randomizes the decision of whether or not there is trade as follows: (i) if $b_1 = b_2 = 0$, there is trade with probability $2 \min\{\Delta, 1/2\}$; (ii) otherwise, there is trade with probability one.
- Conditional on trade, the dealer who submitted the lowest quote sells the asset, and ties are broken uniformly at random.
- Payments are as follows: (i) if $b_1 = -1$ and $b_2 \in \{\bar{b}, 0\}$, trade happens at a price of $-1 + \min\{\Delta, 1/2\}$; (ii) if $b_1 = \underline{b}$ and $b_2 = \bar{b}$, trade happens at a price of \underline{b} ; (iii) otherwise, trade happens at a price equal to b_2 .
- There are no payments made to the platform, so the mechanism balances the budget.
- There is no post-trade transparency. At the end of each period, dealers observe each other's quotes.

Proposition OA.3. *In any equilibrium of the second-price auction with modified payments and rationing, expected welfare coincides with that in the optimal mechanism. Moreover, an ex-post equilibrium exists, so optimal welfare is implementable in the ex-post sense.*

Proof. Step 1: Uniqueness of the equilibrium total probability of trade. We first show that, in any equilibrium, the ex-post *total* probability of trade of the optimal mechanism in Proposition 12 is implemented. In some cases, the identity of the dealer who trades in equilibrium is not uniquely pinned down, and hence individual probabilities of trade need not coincide with what is prescribed by the function $x(\theta_i, \theta_{-i})$. However, the equilibrium outcome will always lead to the same total probability of trade $x(\theta_A, \theta_B) + x(\theta_B, \theta_A)$.

In order to show uniqueness of the equilibrium probability of trade, it is enough to show that: (i) a pair of quotes $(0, 0)$ is observed in equilibrium with probability one if both dealers

are uninformed, and it is observed with probability zero otherwise; and (ii) in any equilibrium and for any pair of quotes on the equilibrium path, the customer must agree to trade with probability one. Indeed, once we ensure that the customer always accepts and that the event $b_1 = b_2 = 0$ is equivalent to the event $\theta_A = \theta_B = 0$, then, by construction of the platform trading rules, the total allocation from the optimal mechanism is implemented.

1.a) Continuation equilibrium in period one. Let b_θ be a quote in the support of the bidding strategy of type θ in period zero, and let us derive equilibrium continuation strategies in period one. Since dealers observe each other's quotes in period zero, both dealers will share a common belief about the asset's value in period one, which we denote by $\hat{v} \in \{-1, 0, 1\}$.

If $\hat{v} = 1$, there are multiple equilibrium strategies in period one; e.g., both dealers may quote 1, or the originally informed dealer may quote 1 and the originally uninformed dealer may quote \bar{b} or vice versa. Nevertheless, all possible equilibria must lead to trade with probability one at a price of 1. This is because the customer in period one will learn that the state is high by observing dealers' quotes and will strictly prefer to trade at any price below $1 + \Delta$. Clearly, trade cannot happen at a price lower than 1 in equilibrium, so it must be that at least one of the dealers quotes 1.

If $\hat{v} = 0$, both dealers are uninformed. Since equilibrium is separating and we already argued that at least one dealer quotes 1 when $\hat{v} = 1$, it cannot be that dealers quote 1 at this node of the game. Since quotes below 0 lead to strictly negative payoffs, it must be that they both quote 0.

Finally, consider the case $\hat{v} = -1$. The only quotes that can be played after this history while still preserving separating strategies for dealers are in $\{-1, \underline{b}, \bar{b}\}$. If both dealers were initially informed, symmetry requires that they both provide the same quote. It cannot be that they both bid \bar{b} , for if they did the customer would refuse to trade, and then dealers would profit from deviating to \underline{b} which would lead to trade at a profitable quote. If they both bid \underline{b} , then the customer would agree to trade with probability one, and dealers could profitably undercut each other by deviating to -1 . Hence, it must be that they both quote -1 . If one dealer was initially informed and the other one was uninformed, there are three kinds of continuation equilibria that are consistent with our equilibrium refinements (although only one of them will survive in the end): (i) both dealers quote -1 , (ii) the type -1 dealer quotes -1 and the type 0 dealer quotes \underline{b} , and (iii) the type -1 dealer quotes -1 and the type 0 dealer quotes \bar{b} . Note that the uninformed dealer gets a zero payoff in all three cases, whereas the informed dealer gets $\underline{b} + 1$ if the second-type equilibrium is played, gets at most $\min\{\Delta, 1/2\}$ if the third-type equilibrium is played, and gets 0 otherwise.

1.b) Equilibrium strategies in period zero. First, it must be that type 1 quotes 1 with probability one. This is because, if she quotes a price less than 1, she trades at a loss in

period zero with strictly positive probability (by symmetry) and still makes at most zero profits in the following period. Next, we look at the uninformed and the low types. For strategies to be separating, it must be that their quoting strategies are supported on some subset of $\{-1, \underline{b}, \bar{b}, 0\}$. Also, note that the continuation equilibrium is unique and yields a zero payoff in all cases, except when quotes reveal that one dealer has a low type and the other one is uninformed, in which case there are three candidate equilibria described previously. For $j = 1, 2, 3$, let $\rho_j \in [0, 1]$ be the probability that players coordinate on the j th-type equilibrium following this history. Note that our restriction to stationary strategies requires that (ρ_1, ρ_2, ρ_3) do not depend on the actual quotes that are given in period t as long as they induce the same hierarchy of beliefs over the dealers' types. For $\theta \in \{-1, 0\}$, let \underline{b}_θ and \bar{b}_θ be the minimum and the maximum quote in the support of type θ 's strategy, respectively, and let $F_\theta(b)$ be the corresponding cdf, and $p_\theta(b)$ the probability mass function (which exists because the auction has a finite bid space).

First, we show that $\underline{b}_0 > -1$. This is because, if $\underline{b}_0 = -1$, the uninformed dealer's belief conditional on winning the auction is strictly higher than -1 (by symmetry), and hence this quote would lead to a negative profit. Second, we show that it must be that $\bar{b}_{-1} < \underline{b}_0$. Suppose otherwise. The previous step implies that $\underline{b}_0 \in \{\underline{b}, \bar{b}\}$, which implies that the low-type dealer does not trade in period t when she quotes $\bar{b}_{-1} \in \{\bar{b}, 0\}$ (this is because the customer will always refuse to trade whenever the winning bid is \bar{b}_{-1}), and hence her overall payoff from this quote is equal to her continuation payoff, which is $\delta\gamma(1-q)(\rho_2(\underline{b}+1) + \rho_3\lambda_1 \min\{\Delta, 1/2\})$, where λ_1 is the probability that the buyer agrees to trade when $Z_t = 1$ and she observes a pair of quotes equal to $(-1, -1 + \min\{\Delta, 1/2\})$ (which need not be equal to 1 given that the buyer is indifferent when $\Delta \leq 1/2$). By deviating to $\underline{b}_0 \in \{\underline{b}, \bar{b}\}$ she gets to trade with strictly positive probability at a price higher than -1 in the current period in the event that the opponent is uninformed. Moreover, if she quotes 0 in the following period after learning that the opponent is uninformed, she can ensure a continuation payoff of $1/2$. Overall, this deviation yields a strictly higher payoff in the current period and a weakly higher continuation payoff, a contradiction.

Next, we show that $\underline{b}_0 = 0$. Suppose that $\underline{b}_0 \in \{\underline{b}, \bar{b}\}$. This implies by the previous step that the low type quotes -1 with probability one. The payoff from this quote is

$$(1-q)\frac{F_0(\underline{b}_0)}{2}\underline{b}_0 < 0,$$

which is strictly negative, a contradiction. Hence, $\underline{b}_0 = \bar{b}_0 = 0$. This in turn implies that, in period zero, the platform will randomize if and only if the two dealers are uninformed.

Finally, we show that $\underline{b}_{-1} = \bar{b}_{-1} = -1$. First, we note that $\bar{b}_{-1} < \bar{b}$. Otherwise, by quoting

\bar{b} , the low-type dealer would not trade (either because she loses the auction or because the customer refuses to trade), and could do strictly better by deviating to -1 . Next, we rule out $\bar{b}_{-1} = \underline{b}$. To do this, we write the low type's payoff from quoting -1 , which is

$$qp_{-1}(\underline{b})(\underline{b} + 1) + (1 - q)\lambda_0 \min\{\Delta, 1/2\} + \delta\gamma(1 - q)(\rho_2(\underline{b} + 1) + \rho_3\lambda_1 \min\{\Delta, 1/2\}),$$

where λ_0 is the probability that the buyer with belief -1 agrees to trade in period zero when the platform prescribes trading at price $-1 + \min\{\Delta, 1/2\}$. When she quotes \underline{b} , she gets

$$\frac{q}{2}p_{-1}(\underline{b})(\underline{b} + 1) + \delta\gamma(1 - q)(\rho_2(\underline{b} + 1) + \rho_3\lambda_1 \min\{\Delta, 1/2\}),$$

where we used the fact that the buyer will refuse to trade at a price greater than $-1 + \Delta$, which is what the platform prescribes whenever i quotes \underline{b} , and the opponent is uninformed and bids 0. If $p_{-1}(\underline{b}) > 0$, then the payoff from quoting -1 is strictly higher, a contradiction.

Observe that, under the above quoting strategies for the dealers, the buyer will agree to trade with probability one, except in the event that $\Delta \leq 1/2$ and the platform prescribes trading at a price of $-1 + \Delta$ (which only happens when one dealer has type -1 and the other one has type 0) in which case the buyer is indifferent. In particular, this says that when $\Delta > 1/2$ the efficient outcome is uniquely implemented.

For the case $\Delta \leq 1/2$, we can use the type -1 dealer's incentive-compatibility constraint to show that the buyer must in fact agree to trade with probability one when she is indifferent, and that the continuation equilibrium when dealers' types are $(-1, 0)$ must be of type (iii). To see this, we write the condition ensuring that the type -1 dealer does not want to deviate and quote 0 in both periods when $\Delta \leq 1/2$, which is

$$(1 - q)\lambda_0\Delta + \delta\gamma(1 - q)(\rho_2(\underline{b} + 1) + \rho_3\lambda_1\Delta) \geq (1 - q)\Delta + \delta\gamma(1 - q)\Delta.$$

The above can be satisfied if and only if $\lambda_0 = \lambda_1 = \rho_3 = 1$. Once we establish that the customer will accept to trade after any on-path pair of quotes and that the continuation equilibrium after a pair of quotes $(-1, 0)$ must be the one where the low-type dealer trades for sure, it follows immediately that the total trade probability resulting in any equilibrium is that from the optimal mechanism.

Step 2: Construction of an ex-post equilibrium. Consider the following quoting strategies.

In period zero, a dealer of type θ gives a quote equal to θ . In period one, if dealer i was initially uninformed and her opponent quoted -1 in period zero, then dealer i quotes \bar{b} ; if i was initially uninformed and her opponent quoted 1 in period zero, she quotes \underline{b} ; in all other cases, dealers quote their initial type θ regardless of what they observed in the previous

period. The customer agrees to trade with probability one whenever she is indifferent.

Let us check the ex-post incentive constraints. First, they are clearly satisfied at all histories for a dealer of type 1, since quoting 1 is a weakly dominant strategy for her. Second, let us look at the uninformed dealer. In period one, if she learned that the opponent has the high type, by quoting \underline{b} , dealer i gets to trade with probability one at a price of 1 and hence gets a zero profit, which is also what she gets if she deviates to any other quote. The same is true after learning that the opponent is uninformed. After learning that the opponent had type -1 , by quoting \bar{b} , she trades with zero probability. There is no profitable deviation here since, by lowering her quote, she only gets to trade at a price equal to the asset's value. Next, consider the uninformed dealer in period zero. If the opponent has type 1, the uninformed dealer may only trade at a price equal to 1 and get a continuation payoff of 0 in period one regardless of what quote she provides, and hence quoting 0 is optimal. If the opponent is uninformed, by quoting 0, dealer i gets a zero payoff in period zero and in the following period. If she deviates to quoting 1, dealer i will trade with probability zero in period zero and get a zero payoff in the following period, given that her uninformed opponent's continuation strategy prescribes quoting $\underline{b} < 0$ after observing the opponent quoted 1. Hence, this deviation is not profitable. If she deviates to quoting $b \in \{\underline{b}, \bar{b}\}$, she can at best trade at a price of 0. Moreover, if the opponent's beliefs assign probability one to this deviation coming from the low type, then this deviation leads to at most a zero payoff in the following period. Similarly, by deviating to -1 , she trades at a price equal to her estimate of the asset's value in period zero and cannot make a positive profit in the following period since she will be perceived as the low type and the opponent will provide a negative quote. Finally, if the opponent has the low type, any quote different than -1 leads to trading with probability zero in period zero and a zero payoff in period one. If dealer i deviates to -1 , she gets to trade with positive probability in both periods at a price of -1 , and hence this deviation yields a zero payoff as well.

Third, consider dealer i with type -1 . If the opponent is uninformed, then dealer i trades with probability one at a price of $-1 + \min\{\Delta, 1/2\}$ in both periods. If she deviates to $b \in \{\underline{b}, \bar{b}\}$ and the customer's and i 's opponent's off-path beliefs assign probability one to this deviation coming from the low type, then this deviation leads to no trade in period zero (since the customer will reject trading at a price of 0) and at most a payoff of $\min\{\Delta, 1/2\}$ in the following period. This is overall strictly worse than what dealer i gets from quoting -1 . If dealer i deviates to 0, she gets to trade at a price of 0 in the current period with probability $\min\{\Delta, 1/2\}$. In the following period, the opponent will believe that i 's type is 0 and will quote 0 again. Dealer i can then either quote 0 as well and trade at this price with probability $\min\{\Delta, 1/2\}$, or quote -1 and trade with probability one at a price

of $-1 + \min\{\Delta, 1/2\}$, both alternatives yielding the same payoff. Hence, her payoff from deviating to 0 is $(1 + \delta\gamma) \min\{\Delta, 1/2\}$, which is exactly her equilibrium payoff. If dealer i deviates to 1, she will get a payoff of 0 in the current period. In the following period, the uninformed opponent will believe that the state is high for sure, and her continuation strategy prescribes quoting \underline{b} . Given this, the most dealer i can obtain in the following period is $\underline{b} + 1 < \min\{\Delta, 1/2\}$. In sum, this deviation yields a strictly smaller payoff compared to i 's equilibrium payoff, so it is unprofitable. Finally, if the opponent has a low type, dealer i gets zero payoff in equilibrium, which is the same as she would obtain for any deviation.

Under all pairs of equilibrium quotes described above the customer is willing to trade, thus confirming that it is optimal for her to always agree to trade on equilibrium path; hence, the outcome implemented in this equilibrium is the same as the one in Proposition 12. \square